



Early Detection of Multiwavelength Blazar Variability

Hermann Stolte¹ , Jonas Sinapius² , Iftach Sadeh² , Elisa Pueschel^{2,3} , Matthias Weidlich¹ , and David Berge^{1,2}

¹ Humboldt Universität zu Berlin, Unter den Linden 6, 10117 Berlin, Germany; hermann.stolte@hu-berlin.de

² Deutsches Elektronen-Synchrotron DESY, Platanenallee 6, 15738 Zeuthen, Germany; jonas.sinapius@desy.de

³ Ruhr-Universität Bochum, D-44780 Bochum, Germany

Received 2024 October 15; accepted 2024 November 20; published 2025 February 7

Abstract

Blazars are a subclass of active galactic nuclei with relativistic jets pointing toward the observer. They are notable for their flux variability at all observed wavelengths and timescales. Together with simultaneous measurements at lower energies, the very-high-energy (VHE) emission observed during blazar flares may be used to probe the population of accelerated particles. However, optimally triggering observations of blazar high states can be challenging. Notable examples include identifying a flaring episode in real time and predicting VHE flaring activity based on lower-energy observables. For this purpose, we have developed a novel deep learning analysis framework, based on data-driven anomaly detection techniques. It is capable of detecting various types of anomalies in real-world, multiwavelength light curves, ranging from clear high states to subtle correlations across bands. Based on unsupervised anomaly detection and clustering methods, we differentiate source variability from noisy background activity, without the need for a labeled training data set of flaring states. The framework incorporates measurement uncertainties and is robust given data quality challenges, such as varying cadences and observational gaps. We evaluate our approach using both historical data and simulations of blazar light curves in two energy bands, corresponding to sources observable with the Fermi Large Area Telescope and the upcoming Cherenkov Telescope Array Observatory. In a statistical analysis, we show that our framework can reliably detect known historical flares.

Unified Astronomy Thesaurus concepts: [Blazars \(164\)](#); [Transient detection \(1957\)](#); [Outlier detection \(1934\)](#); [Neural networks \(1933\)](#); [Gamma-ray astronomy \(628\)](#); [Astronomical methods \(1043\)](#); [Extragalactic astronomy \(506\)](#)

1. Introduction

Observations of transient phenomena are key to disentangling the physical processes at play in astrophysical systems. In past years, notable transient events have included the spatial correlation of an astrophysical neutrino with a flaring gamma-ray source (M. G. Aartsen et al. 2018), the observation of gamma rays and gravitational waves produced by a kilonova (B. P. Abbott et al. 2017), and the discovery of new classes of very-high-energy (VHE; >100 GeV) gamma-ray emitters such as gamma-ray bursts and novae (H. E. S. S. Collaboration et al. 2019; MAGIC Collaboration et al. 2019; F. Aharonian et al. 2022) via detection of transient emission. Measuring or constraining the VHE gamma-ray emission from transient events is of particular interest, as these gamma rays track the most extreme acceleration processes.

Imaging atmospheric Cherenkov telescopes (IACTs), such as VERITAS, MAGIC, H.E.S.S., and the next generation Cherenkov Telescope Array Observatory (CTAO), are the most sensitive instruments for measuring VHE gamma-ray emission (F. Aharonian et al. 2006; B. S. Acharya et al. 2013; N. Park 2015; J. Aleksić et al. 2016). However, they have fields of view of less than 10° , which limits the chance of serendipitous transient detection. Source variability detected at other wavelengths is therefore used to trigger IACT observations, increasing the probability of observing transient events with IACTs. The deep learning approach presented here collates multiwavelength information on the activity state of

known gamma-ray sources and uses this information to anticipate periods of unusual emission (e.g., flares) in the VHE band. In developing this method, we focus on blazars, a gamma-ray source class which shows strong variability on all observed timescales and wavelengths, and for which a large archive of VHE and multiwavelength observations exists. The code used in this work is available on GitLab at https://gitlab.desy.de/trans_finder/blazar_flares.

1.1. Blazars

Blazars are a class of active galactic nuclei with a relativistic jet oriented toward the observer, resulting in strongly Doppler-boosted emission. They are the most numerous source class detected in high-energy (HE; 100 MeV–100 GeV) and VHE gamma rays. However, the mechanisms driving particle acceleration in the jets and the observed gamma-ray emission are not well understood, and models of varying complexity abound (for discussion of acceleration, see, e.g., J. G. Kirk et al. 2000 and L. Sironi et al. 2015; for emission, see M. Cerruti 2020, and references therein).

Gaining a better understanding of the underlying acceleration and emission mechanisms in blazars is relevant for a number of topics. For example, blazars are plausible sources of ultra-high-energy cosmic rays, should protons be accelerated in their jets (e.g., K. Murase et al. 2012; X. Rodrigues et al. 2018). As extragalactic sources, located at cosmological distances from Earth, blazar observations can be used to probe the photon and magnetic fields traversed en route to Earth, and to test for signatures of Lorentz invariance violation and photon coupling to axion-like particles, effects that are expected to grow with gamma-ray propagation distance (see J. Biteau & M. Meyer 2022, and references therein).



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

For such scientific goals, precise measurements of the broadband spectral energy distributions (SEDs) of sources is critical, being the main handle for interpretation of the emission mechanisms.

The photon emission from all VHE-detected blazars is characterized as variable at some or all observed wavelengths, from radio to VHE, on timescales from minutes to years. Bright states are of particular interest, but it is challenging to simultaneously detect a source and measure the SED over multiple wavelengths. This is due in part to limitations on the collection areas of instruments in different wavelengths, as well as to the intrinsic variation in flux as a function of energy. Sources that are typically too dim to be detected in the HE and VHE bands without long integration periods (weeks to months) can conversely be detected during flaring episodes within minutes or hours.

The SED of a variable blazar is modeled most robustly using simultaneous multiwavelength observations, in order to ensure that different time-dependent flux states are not conflated. Bright states are key to probing the emission mechanism of blazars, and to their physical interpretation. Particularly important is the VHE band, where the photon flux is comparatively low. Flare detection is therefore a cornerstone of the scientific programs of the IACT community (see F. Aharonian et al. 2007; J. Albert et al. 2008; V. A. Acciari et al. 2009; A. U. Abeysekara et al. 2015, for selected examples). Identifying correlated variation at different wavelengths is also valuable, as different models of blazar emission predict correlations (or lack thereof) between different bands. However, consistent definition and identification of flares remains a challenge in the community, as discussed for example by R. Zimmerman et al. (2024).

1.2. Deep Learning and Anomaly Detection

Machine learning, in particular deep learning, are widely used in astronomy. Many applications are based on supervised learning. This involves input data which are labeled, where the algorithm is trained to encode the mapping between inputs and labels. Examples include classification, such as γ /hadron separation for IACT experiments (see, e.g., Q. Feng & T. T. Y. Lin 2016; D. Nieto Castaño et al. 2017), as well as regression tasks, e.g., the evaluation of the redshift of a galaxy (I. Sadeh et al. 2016). Conversely, for unsupervised learning, specific labels are not known a priori and the objective is to find patterns in the data. Applications are commonly based on clustering (E. Min et al. 2018) and/or outlier detection (I. Reis et al. 2018; K. L. Malanchev et al. 2021).

In the following, we expand on the work of I. Sadeh (2020, hereafter SA20), who utilized a recurrent neural network (RNN) for outlier detection. An RNN is a form of directed graph, representing a sequence of steps in time. Outputs from each time step are fed as input to the next, in addition to the respective temporal data. Such models are useful for characterization of complex data on different timescales.

SA20 illustrated the use of their method for the detection of various types of astrophysical transients, such as gamma-ray bursts and neutrino emission by candidate neutrino point sources. The general concept is to use an RNN to characterize the background to a potential transient. This is done by providing the network with existing/past data, taken before the emergence of the putative transient event. For instance, in the case of the search for an astrophysical neutrino transient, the

background represents the continuously observed signals of atmospheric neutrinos. The purpose of the RNN is to predict the background within the near future (upcoming time steps of the network) based on observations of the near past. Potential transients are then detectable as deviations of new observations from the expected background predictions.

For their chosen data sets, SA20 could construct RNN inputs that increase with the intensity of putative signals. For example, they used the reconstructed gamma-ray flux within a region of interest (ROI) from an IACT experiment, which increases when a new gamma-ray source appears. As such, transients would always manifest themselves as upward fluctuations from the background predictions of the network. Considering observables that by construction always scale up with the strength of the signal enables the definition of a simple outlier score, acting as a test statistic (TS) for detection. The latter is defined by SA20 as the integrated difference between the predictions of the RNN and the stream of real data.

In the current work, we focus on blazar flares as the target transient phenomena. However, we set out to identify any unexpected activity from the source, such as notable downward deviations or correlated changes of the emission on different timescales. For this purpose, the simple TS used by SA20 is not appropriate, given that it is designed exclusively for upward deviations from the background. In order to construct a generalized TS for arbitrary types of flares, we expand the architecture of the network used by SA20, as discussed in the following.

This paper is organized as follows. In Section 2, we introduce our machine-learning-based framework, and in Section 3 we describe how it is trained. Section 4 contains a simulation study, in which we evaluate the sensitivity of our method. In Section 5, we demonstrate the utility of our method in an analysis of real data from BL Lacertae, before concluding in Section 6.

2. Anomaly Detection Pipeline

Our pipeline is illustrated in Figure 1 and the most important concepts and parameters are summarized in Table 1. After preprocessing of the input light curves, the data are organized into a time-series format. We distinguish between a *context window* and a *search window* for the analysis. In the latter, we scan for putative flares.

Forecasting (Box I in Figure 1). The inputs within the context window, denoted by τ_{enc} , and the search window, τ_{dec} , are respectively mapped into the encoder and decoder components of an RNN, similar to the approach of SA20. The purpose of the encoder is to provide contextual information (recent source activity) for the decoder. We train the model under the assumption that the context window is devoid of potential signals, and thus represents a realistic background for flares. The outputs of the decoder, τ_{for} , are the model predictions for the light curve, given this background hypothesis.

Embedding (Box II). The predictions are contrasted with the actual data in order to quantify anomalies. As explained in Section 2.3, we derive the difference between the two sets as $\tau_{\text{res}} = \tau_{\text{dec}} - \tau_{\text{for}}$, emphasizing the most recent data in time through a weighting scheme. We model the distribution of τ_{res} for the background using an autoencoder, which compresses the data into a low-dimensional space. For the subset of data which comprises background configurations, τ_{res} represents the

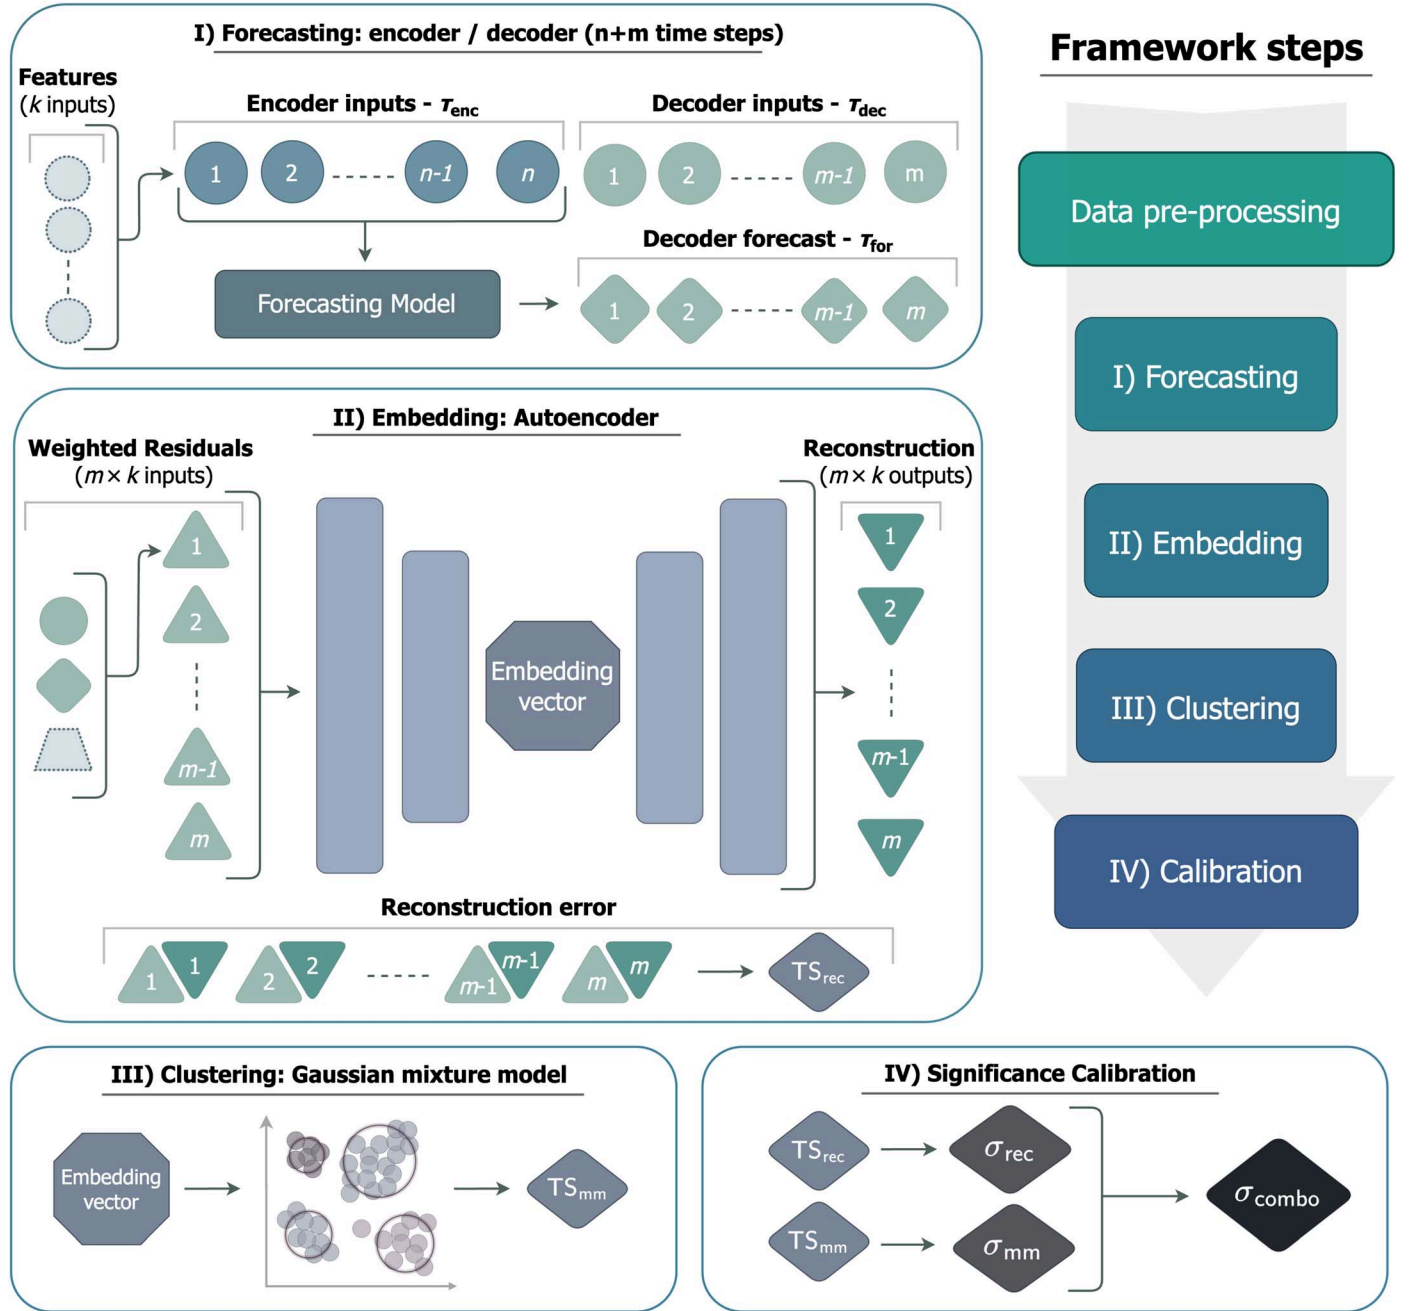


Figure 1. Illustration of the architecture of the model, as described in the text.

intrinsic variations of the light curve in the absence of flares. Otherwise, τ_{res} quantifies the strength of a potential anomaly.

Clustering (Box III). The condensed representation provided by the autoencoder facilitates comparison between new data and the reference distribution, which comprises all known background states. We parameterize the distribution of τ_{res} for the background data set with a Gaussian mixture model (a form of clustering), as discussed in Section 2.4.

Significance calibration (Box IV). The compatibility of new data with the modeled distribution of τ_{res} is used for anomaly detection. We follow the approach of SA20 in order to map the corresponding TS into a significance for detection. A detailed description and derivation of this calibration is given in the Appendix of SA20. In addition, we derive another complementary TS directly from the outputs of the autoencoder, as

discussed below. The final p -value for flare detection is derived from the combination of these two statistics (Section 2.5).

Once a significant flare is detected, it raises the question of which elements in the data are driving the detection. We analyze the contribution of different data to the overall significance within the pipeline, as described in Section 2.6. Below, the various elements of the pipeline are discussed in detail.

2.1. Light Curves as Input Data

The multiwavelength inputs to our pipeline are a set of light curves obtained from different instruments. Each input (which we refer to in the following as a *channel*) captures an

Table 1
Overview of Notions and Parameters

Notation	Description
t_{ref}	Reference time pivot in the input light curve
$\tau_{\text{enc}}, \tau_{\text{dec}}$	Inputs split into encoder (context) and decoder (search) windows, comprising n and m time steps
τ_{for}	Forecast of background activity for m time steps
$\tau_{\text{res}}, \tau'_{\text{res}}$	Residuals between decoder, τ_{dec} , and forecast, τ_{for} ; tagged variable is adjusted by temporal weights
ω	Temporal weights applied to decoder residuals
τ_{rec}	Reconstruction of the weighted residuals, τ'_{res}
ϕ	Embedding vector representing residuals, τ'_{res}
$\text{TS}_{\text{rec}}, \sigma_{\text{rec}}$	Test statistic (TS) and significance, based on the reconstruction error
$\text{TS}_{\text{mm}}, \sigma_{\text{mm}}$	TS and significance, based on the clustering model compatibility metric
σ_{combo}	Combined significance from TS_{mm} and TS_{rec}
γ_{timeaggr}	Time bin width (in days for the current study) to be considered a single time step
γ_{signoise}	Threshold (in units of standard deviation) to constrain variability; used to quantify background states for training
γ_{decay}	Temporal decay rate, used for suppressing old anomalous data

observable related to source brightness (e.g., flux, magnitude), with associated uncertainties.

A potential flare is always searched for within a specific temporal context. The data collection process behind each channel, however, is typically affected by numerous observational constraints, such as the visibility of the source and weather constraints for ground-based instruments. Consequently, the observing cadences commonly vary across channels and over time. Frequently, limited data availability and long gaps in coverage need to be handled. To address these temporal dynamics, we employ a series of techniques, which we discuss in the following.

Consistent time binning. To jointly analyze observables in the timescales that are relevant for blazar flares, we first ensure a consistent time binning in each channel. Specifically, we contrast between high-density and low-density channels, given the selected width of our time binning, γ_{timeaggr} (e.g., one day, customizable per channel). High-density channels, which generally contain multiple data points within γ_{timeaggr} , are downsampled by randomly selecting a single data point per γ_{timeaggr} . This effectively captures the average state in the channel.

Dynamic windowing. A potential flare is searched for in the data that is available up to a time pivot, t_{ref} . Observation cadences and the corresponding coverage will generally vary across inputs and over time. The number of data points within a fixed time window may therefore differ between channels. To effectively interpret source activity in context, it may be necessary to consider older data from sparse channels together with recent data from dense channels. Instead of relying on fixed-duration windows, it is hence desirable to consider time windows of variable lengths. We employ a dynamic-windowing approach, which considers a fixed number of the most recent available data from each channel, up to t_{ref} . The first (i.e., oldest) n data points in the window constitute the context window for flare detection, τ_{enc} ; the last (i.e., newest) m data points form the search window, τ_{dec} .

Variable and sparse coverage. In order to constrain emission models and predict future flares, the observation time of an individual measurement is crucial. However, the provenance of the available multiwavelength data is subject to potentially arbitrary constraints. The cadence of measurements is influenced by several confounding factors, including observing conditions, source visibility, and instrumental availability. Correspondingly, the presence of a particular measurement could itself be seen as abnormal. While our framework is designed to consider any type of anomaly in the input channels, it must be agnostic to changes in data cadence, while modeling the relevant temporal context. For this reason, we consider the difference in cadence between measurements only indirectly through a dynamic weighting scheme, as described in the following.

In the assessment of potential flaring activity in channels with varying sparsity, recent data are generally considered more relevant than older data. This is particularly important for low-density channels, which generally contain a single data point within a period of several γ_{timeaggr} . The search window of a low-density channel may consist entirely of data that lie far in the past (relative to t_{ref}). It may thus capture flaring behavior that is irrelevant for assessing the current activity of the source.

We use a scheme of temporal weights to ensure the sensitivity to potential new flares, given the possible presence of old anomalies. Weights are defined as

$$\omega(t_{\text{step}}, \gamma_{\text{decay}}) = (1 - \max(0, t_{\text{step}} - t_{\text{delay}}))^{-\gamma_{\text{decay}}}. \quad (1)$$

Here, t_{step} is the temporal distance of a particular data point relative to t_{ref} . We define t_{delay} as the period during which all data are considered recent. This parameter is defined by the number of decoder steps, m , and the time step interval, γ_{timeaggr} , as

$$t_{\text{delay}} = m \times \gamma_{\text{timeaggr}}. \quad (2)$$

The parameter, γ_{decay} , controls the strength of the temporal decay.

The temporal weights are used to scale the inputs to the autoencoder. Weights are constructed such that they are equal to unity in search windows that are not sparse. This scheme serves to relatively shift the residuals for older data toward background-level values. Thus, the impact of anomalies which are driven by source variability in the far past is suppressed.

The temporal weighting scheme is illustrated in Figure 2. Here, the top panel shows an example input light curve comprising two channels and a time pivot, t_{ref} . The first channel has dense coverage and thus a short window size. The second channel is more sparsely sampled. The corresponding window covers a wider time interval, as required to capture $n + m$ data points. The lower panel shows the corresponding time-series structure separated into the context window, τ_{enc} , and the search window, τ_{dec} . The dynamic weighting scheme and the effect of γ_{decay} is also illustrated in the two bottom panels. Here, the data in the sparse channel are shifted toward background within the decoder, as there exists a notable gap between the most recent data and t_{ref} . In contrast, the weights of the dense channel leave the data within the decoder unmodified.

Limited data availability. The availability of data in a channel may be limited, for instance due to novel instruments starting operations, or to having made only few observations of a source. In some cases, a channel may contain less data than required for the construction of a single window. To ensure the

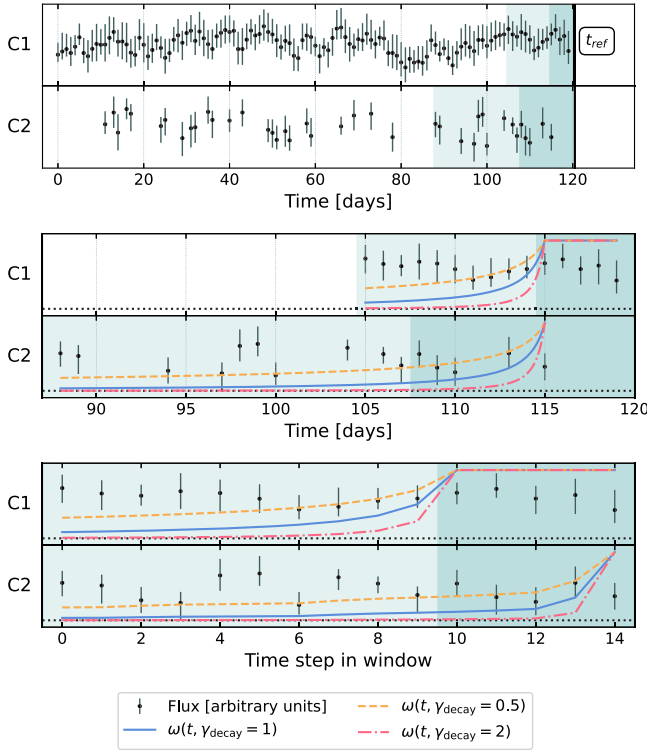


Figure 2. Illustration of a dynamic-window transformation of two light curves into a combined time-series structure. The top panel shows two light-curve channels, C1 and C2, with a reference time pivot, t_{ref} . The context window, τ_{enc} , and the search window, τ_{dec} , are highlighted as light and dark shaded regions. The middle panel shows a zoom in of both windows over the time relative to t_{ref} . The bottom panel shows both windows in $n + m$ discrete time steps as they are ingested into the forecaster and autoencoder. Colored lines in the lower panels show the corresponding temporal weights, ω , for three representative settings of γ_{decay} . The dotted horizontal line marks a weight, ω , of zero.

compatibility of our approach with limited data, we add padding to the beginning of the window. It consists of new data points, derived from existing data, randomly scrambled in time. This includes time stamps lying in the far past. With the aforementioned weighting scheme, we ensure that the padding segment of data does not result in anomalies.

Uncertainty on observables. In order to account for the uncertainties on individual measurements, we draw a sample of γ_{err} time series with feature values sampled from the assumed normally distributed, potentially asymmetric uncertainties of the data. While this is a common approach, the framework remains agnostic to the choice of distribution used for modeling uncertainties on observables. The following stages independently process the original time series as well as all uncertainty samples. In the final calibration, as discussed below, their corresponding TSs are aggregated. The final outcome is then a calibrated significance with an associated uncertainty which represents the propagated input uncertainties from all channels.

2.2. Forecasting Background Activity

After preprocessing, the input light curves are transformed into fixed-size windows that follow a time-series structure of $n + m$ time steps. Each step contains k features, one for each input channel. The specific values chosen for these parameters are specified in Section 4.2. As illustrated in Figure 1 (Box I), a window is then split into two parts, an encoder, τ_{enc} , and a

decoder, τ_{dec} . The encoder serves as the reference for a forecast, τ_{for} , representing the predicted background-level source activity. Both τ_{dec} and τ_{for} comprise the same m time steps following the n time steps in τ_{enc} . To obtain the forecast of background activity, we employ a multivariate, multistep RNN with probabilistic output layers, similar to the encoder-decoder architecture utilized by SA20.

2.3. Representation Learning

After the forecasting of background activity, we contrast the actual data with the background forecast, τ_{for} . That is, we derive element-wise residuals as

$$\tau_{\text{res}} = \tau_{\text{dec}} - \tau_{\text{for}}. \quad (3)$$

Each element in τ_{res} is then multiplied with its corresponding temporal weight,

$$\tau'_{\text{res}} = \omega \cdot \tau_{\text{res}}, \quad (4)$$

where the elements of τ'_{res} serve as the inputs to the next element of the model.

This definition of τ'_{res} is motivated by the fact that small numerical values of the residuals correspond to background states, while larger positive or negative values represent increasingly significant anomalies. Effectively, this weighting scheme shifts older data toward background states, with $\omega = 1$ corresponding to the original data and $\omega = 0$ being fully compatible with background.

Based on the weighted residuals, τ'_{res} , we obtain an embedding vector, ϕ , that represents the source variability across channels. The purpose of the mapping into an embedding vector is to capture a condensed representation of variability across all channels. It enables the comparison of relative distances between source activity states, which is fundamental to one of our two test statistics, TS_{mm} , used for flare detection (see Section 2.4).

To map τ'_{res} into an embedding vector, we employ a variational autoencoder with probabilistic output layers based on RNNs. As the autoencoder is a modeled distribution of observed source activity, it can fail to generalize to novel source activity. This could happen if the novel activity is too different from the data used for training, and would result in unreliable mappings into the embedding space. We interpret such a case as a potential type of flaring state, where the differences to known source activity are so pronounced that a comparison in the embedding space is infeasible.

To reliably detect this type of flare, we compare the reconstruction outputs from the autoencoder, τ_{rec} , to the inputs, τ'_{res} . When the embedding vector is unreliable, the corresponding reconstruction error is high. Hence, we define a test statistic, TS_{rec} , as the compatibility of each output with its corresponding input distribution. Specifically, this TS is calculated as the summed negative log-likelihood of each element in τ'_{res} with respect to the corresponding distribution reconstructed from the embedding vector:

$$\text{TS}_{\text{rec}}(\tau'_{\text{res}}, \tau_{\text{rec}}) = \sum_{i=0}^m \sum_{j=0}^k -\log p(\tau_{\text{rec},i,j} | \tau'_{\text{res},i,j}). \quad (5)$$

The mapping into embedding space and derivation of TS_{rec} is illustrated in Figure 1 (Box II).

2.4. Clustering

The embedding vector, ϕ , obtained from the autoencoder, represents an unknown state of source activity. Next, we calculate a test statistic, TS_{mm} , that quantifies the compatibility of this state with a model for the background. We use a Bayesian Gaussian mixture model (BGMM; G. J. McLachlan & D. Peel 2000; D. M. Blei & M. I. Jordan 2006), χ , for modeling background activity in the embedding space. The goal is to allow for complex clusters of different background states to be jointly modeled.

The calculation is illustrated in Figure 1 (Box III). Intuitively, one can relate TS_{mm} to the geometric distance between the embedding vector of a source state and the distribution of modeled background states derived from the BGMM. The “farther away” the source state embedding lies from the background embeddings, the less likely it is for this source state to be compatible with the background model. Technically, the TS is derived as the negative log probability of the embedding vector with respect to the mixture model:

$$TS_{\text{mm}}(\phi) = -\log p(\chi|\phi). \quad (6)$$

2.5. Significance Derivation

Box IV in Figure 1 illustrates the final step of the pipeline, carried out after the model has been fully trained. Here, we map the two test statistics, TS_{mm} and TS_{rec} , into their respective p -values for anomaly detection, p_{mm}^v and p_{rec}^v . As defined above, the first metric, TS_{mm} , represents in-distribution anomalies (low compatibility with the background model). The second metric, TS_{rec} , captures out-of-distribution anomalies (large reconstruction errors of the autoencoder). The p -values are conveniently represented by their corresponding significance values, the clustering significance, σ_{mm} , and the reconstruction significance, σ_{rec} .

In addition to these two quantities, we derive a high-level TS. This metric is based on the two individual p -values, and is defined as

$$TS_{\text{comb}} = -\log p_{\text{rec}}^v - \log p_{\text{mm}}^v. \quad (7)$$

The combined TS is similarly used to derive the final p -value and significance for the pipeline, p_{comb}^v and σ_{comb} .

The so-called *calibration* process between each TS and the related significance for detection follows the prescription of SA20. We employ a numerical approach that involves evaluating each TS multiple times, given a representative sample of background examples. The corresponding distribution of TSs therefore represents the null hypothesis for anomalies. The TSs are constructed such that high numerical values correspond to larger potential anomalies. The distribution for the background data set may therefore be used to estimate the probability for a high TS value (small p -value), and the related significance.

2.6. Characterization of Signals

When a flare is identified by the framework, retracing what aspects of the data contribute to this identification is nontrivial. However, understanding what part of the overall behavior is considered anomalous is key in furthering our understanding of the nature of blazar flares. An anomaly in the behavior of a single channel can constitute a flare, as well as the combined behavior of multiple channels. For instance, a correlation

between two channels may be anomalous, even in cases where the same data would not individually be flagged as unusual.

To guide the identification of anomalous behavior, we assess the significance when considering only subsets or individual channels. This is accomplished by modifying the temporal weights, ω , on τ'_{res} . Specifically, we set $\omega = 0$ for all channels that are not considered in a given iteration during inference. This is effectively a Boolean mask that sets all data from masked channels to correspond to background. In the evaluation, we demonstrate how this enables the tracing of significant detections back to individual input channels.

We note that these projections in the embedding space are not individually calibrated to p -values. That is, the corresponding significance of individual channels is indicative of their relative contribution to anomalies. However, they should not be interpreted as the final significance for detection.

3. Training the Model

All components of the pipeline are trained on data corresponding to background source activity. The available data, however, typically consist of light curves which include flaring periods or other anomalies. To obtain a background-only version of the data, we rely on a combination of (optional) manual filtering to remove obvious flaring states, as well as several automated cleaning and augmentation procedures. Both methods are described in detail in Section 3.1. They include several randomization procedures for statistically removing potential remaining anomalies from the data used for training (forecasting and clustering stages) as well as for calibration. At this stage, the training data set is normalized, such that each channel (i.e., feature) has a mean of zero and a standard deviation of 1.

The data set is augmented in order to increase the robustness of our model. As described in Section 3.2, we introduce randomly generated long-term states to a subset of training examples. As part of training the autoencoder, we also add randomly generated short-term fluctuations, as described in Section 3.3. This augmentation ensures that the embedding space covers a wide range of source activity. It helps to reduce potential reconstruction errors and improves the generalizability of the pipeline. By construction, these fluctuations deviate from the background distribution we intend to model. We therefore emphasize that they are not used for fitting the BGMM or for the significance calibration process.

The training process proceeds as follows.

Step I. The forecasting model is trained on background-only data using the Adam optimizer (D. P. Kingma & J. Ba 2014) with the summed log probability for the loss, a learning rate of 0.005, a dropout of 10%, L2 regularization of 0.0001, and early stopping.

Step II. The autoencoder is trained analogously to the forecasting model. As a notable difference, the training data set is enriched with injected short-term fluctuations. After training the forecasting model, the inputs to the autoencoder, τ'_{res} , are normalized to have a mean of zero and a standard deviation of 1. The temporal weight decay, γ_{decay} , is set to randomized values within the range $\{1, 2\}$ to further enhance the phase space of the embeddings. The optimizer and regularization techniques are analogous to the training of the forecasting stage.

Step III. The Gaussian mixture model of background-level activity is fitted using a Markov Chain Monte Carlo (MCMC)

approach. Here, we adaptively increase the complexity of both the model and the MCMC parameters until there is no notable improvement in the quality of fit. The dynamic parameters include the number of model components as well as the number of MCMC steps, of burning steps, and of step-size adaptations. The quality of fit is determined through a two-sample Kolmogorov–Smirnov test for equal distributions between the input and the posterior sample distribution. To further improve the convergence stability, multiple MCMC chains are fitted in parallel. The best chain is selected according to the quality of fit and is used for subsequent inference. Through adaptively increasing the complexity of the model and the fit, we ensure convergence once the process terminates regularly, i.e., without reaching a time-out.

3.1. Cleaning Anomalies from the Training Data Set

Available data sets for training typically comprise light curves that include flaring periods or other anomalies. As the input to the training stage of the pipeline should consist only of background examples, such data are cleaned prior to training. Well-defined and extreme flaring intervals are removed by hand. This procedure is optional.

Following manual cleaning, some anomalies may remain in the data, such as weak, short-term fluctuations, or other correlations between channels. To ensure sensitivity to these types of anomalies, the pipeline includes additional statistical randomization and cleaning steps.

As part of automated cleaning, we use a variability threshold (in units of standard deviation), defined by the parameter γ_{signoise} . When a particular window contains data with a local signal-to-noise ratio higher than γ_{signoise} , we exclude it from the training and calibration stages. This ensures that the model becomes sensitive only to deviations stronger than the configured threshold. Tuning this parameter is crucial; increasing the value of γ_{signoise} reduces the rate of false positives, but also decreases the sensitivity to faint flares. On the other hand, lowering the threshold raises the rate of false positives, while improving the sensitivity to faint flares.

We also shuffle data inside each window. This suppresses residual correlated anomalies within the training data set. It is important to avoid introducing a bias to the pipeline by always having the same observation cadence in any one channel. We therefore randomize the time stamps of individual data points as follows. Considering data points i and j taken at times $t_i < t_j$, we reassign the time stamp $t_i \rightarrow t'_i$ to a random value that satisfies the condition $t_i \leq t'_i < t_j$ and conforms to the selected time bin width, γ_{timeaggr} . This is done for each data point in each input channel independently.

Finally, we apply a randomized temporal shift to each light curve individually; this is done globally for the entire channel. Multiple randomized realizations of the data set cover a wide range of relative global shifts between channels. Such transformations suppress possible inter-band correlations.

3.2. Timescales for Anomaly Detection

Blazars are known to exhibit variability on short (minutes), intermediate (days), and long (years) timescales. The examples featured in the current study are focused on detecting medium-timescale behavior, on the order of a few days. We account for very short flares with the choice of $\gamma_{\text{timeaggr}} = 1$ day. It therefore remains important to consider the impact of long-term

variability. Over long enough time periods, on the order of a few years, blazar emission can exhibit trends and correlated behavior of different levels of intensity. We would like to avoid false-positive detections for a medium-timescale search. Such slow-changing trends should therefore be made indistinguishable from background in the training data set.

We augment the training sample accordingly. We introduce randomly generated long-term states in the input light curves. This is done by globally shifting the data up or down by up to 6 standard deviations for a particular realization. Here, all data in a given channel have the same vertical offset applied, where individual channels are treated independently. This augmentation also serves to improve the generalization of the modeled background.

3.3. Embedding Space Augmentation

To allow the embedding space to represent different types and intensities of potential flares, we enrich the data set for training the autoencoder by introducing arbitrary short-term fluctuations. These allow the embedding space to encode distances between different types of potential flaring states. Specifically, the random fluctuations are modeled as step functions and Gaussians with amplitudes randomly sampled within $\{2, 100\}$ (upward) or $\{-2, -10\}$ (downward) local standard deviations. Their duration is within $\{1, 8\}$ days. The temporal onset of a fluctuation is randomly shifted across channels by up to five 5.

4. Simulation Study

We evaluate the performance of our pipeline based on simulations of quiescent and flaring states inspired by the TeV-detected BL Lac object 1ES 1215+303. We simulate and evaluate different flare scenarios, varying in duration and flux. The light curves cover observations from Fermi-LAT and CTAO.

The Large Area Telescope (LAT) is a pair-conversion telescope on board the Fermi spacecraft (W. B. Atwood et al. 2009), covering an energy range from ~ 20 MeV to more than 500 GeV. During standard operations it covers the full sky once every 3 hr.

The upcoming CTAO is a collection of VHE gamma-ray IACTs, with one array on the Southern Hemisphere and one on the Northern Hemisphere (Cherenkov Telescope Array Consortium et al. 2019), covering energies between ~ 20 GeV and 300 TeV. With increased sensitivity compared to current IACTs, it will significantly improve knowledge of the VHE sky.

In Section 4.1, we describe the simulation in more detail, before discussing the training process and pipeline parameters in Section 4.2. The results are shown in Section 4.3.

4.1. Data Set

We simulate the multiband gamma-ray light curve of a quiescent source, mimicking the 2017 steady-state spectrum of 1ES 1215+303, as measured by J. Valverde et al. (2020) with Fermi-LAT and VERITAS. The measured Fermi-LAT log-parabola and the VERITAS power law are jointly well described with a log-parabola that accounts for extragalactic background light (EBL) absorption, shown in Figure 3. Assuming constant source behavior, we use the `gtoobsim` tool from `Fermitools` (version 2.2.0), and the `gammapy` package (v 0.20.1; C. Deil et al. 2017) to produce simulations of 1000 days daily-binned Fermi-




Figure 3. Spectral energy distributions for the simulated data set, inspired by the 2017 steady-state spectrum of the blazar, 1ES 1215+303. Horizontal arrows show the sensitivity regions of Fermi-LAT, VERITAS, and CTAO, respectively. The orange dashed (dashed-dotted) line shows the spectrum given in J. Valverde et al. (2020) for LAT (VERITAS), and the solid green (blue) line shows the combined, EBL-absorbed log-parabola used for the baseline (flaring) simulations.

LAT and CTAO light curves, respectively.⁴ To simulate a realistic IACT observing cadence, we conservatively decrease the sampling of the CTAO light curve to mimic the sampling of real VERITAS observations of this source. Subdividing the light curves into multiple, nonoverlapping energy bins allows us to keep some of the spectral information. The Fermi band is divided into three energy bins (100–669 MeV, 669 MeV–4.48 GeV, and 4.48–300 GeV); the CTAO band is divided into four energy bins (31.6–79.4 GeV, 79.4–199.5 GeV, 199.5 GeV–1.2598 TeV, and 1.2598–12.589 TeV). The inputs and specifics of these simulations are described in more detail in Appendix A.

We create templates to evaluate the performance on simulated flares, as described above. As a proxy for a high state, we use the 2017 flare studied by J. Valverde et al. (2020), denoted as F_{ref} in the following. For comparison, the SEDs of the baseline and the unscaled flare are shown in Figure 3. We simulate observations of a flare with a top-hat profile on six subsequent days with 1%, 10%, 50%, 100%, and 150% of F_{ref} . By applying these templates only partially, we can also simulate flares with durations between 1 and 5 days. To increase the richness of the data set and to better account for random fluctuations in these templates, we create 10 randomized realizations for each template at each intensity. While we are generally interested in sudden high-flux states in blazars, an unexpected decrease in flux can also be anomalous and potentially interesting. To test how the framework reacts to such cases, we create realizations where we subtract these templates from the baseline.

4.2. Configuration Parameters and Training

We utilize a time bin width of $\gamma_{\text{timeaggr}} = 1$ day to transform the data into windows consisting of $m = 10$ context steps and $n = 5$ search steps. To filter nonbackground data, a signal-to-noise threshold of $\gamma_{\text{signoise}} = 5$ is applied. We generate 100 augmented realizations of pure background, and 100 augmented realizations including short-term fluctuations. The temporal

weight decay is chosen as $\gamma_{\text{decay}} = 1$. The training converges, which we confirm by inspecting the loss curves and calibration control plots for the goodness of fit. We also require that the adaptive MCMC fit terminates regularly.

4.3. Results

To study the performance of our pipeline, we first show that flaring and background states are distinguishable by their respective embedding vectors obtained from the autoencoder. We then assess that the detection significance scales as expected with varying flare strengths, considering the expected time coverage within the Fermi and CTAO channels.

4.3.1. Background and Flaring States as Embeddings

We expect background and flaring states to be mapped into distinctive embedding clusters. To confirm this, we visualize the embedding vectors of our simulated data set in Figure 4. In the bottom left, we show the distribution and pair-wise projections for each embedding vector dimension. Background states are shown in blue, and flaring states in orange. For illustration, we also show 2D projections of the embedding vectors using standard dimensionality reduction techniques to visualize complex data, namely principal component analysis (PCA; K. Pearson 1901) on the top, and uniform manifold approximation and projection (UMAP; L. McInnes et al. 2018) on the bottom. Here, the color of each point shows the associated significance, with low-significance data (i.e., data consistent with background states) shown in blue, and high-significance data shown in yellow.

It is apparent that background and flaring states are distinct, where some dimensions are more important for distinguishing between states than others. The linear (PCA) and nonlinear (UMAP) projections into 2D both show a notable difference between flaring and background states. A certain overlap between the two distributions is visible and expected, as some aspects of the data can be consistent between background and (weak) flaring states.

Another observation is that flaring states cover a larger space compared to background states in most dimensions of the embedding. The larger spread in the embedding vectors stems from the larger intrinsic variability of light curves in flaring states. The smaller variations in the background states are consequently mapped to a relatively compact region. This is a desired outcome of the model, illustrating that the representations capture features of variability, which are the basis for anomaly detection.

4.3.2. Significance and Flare Strength

The detection significance is expected to scale with the strength of a flare. To study this, we show the significance as a function of flare strength in the three panels on the left of Figure 5. The leftmost panel shows the combined significance, while the next two panels show the mixture model significance (top) and the significance of the reconstruction error (bottom). In each panel, we show the respective significance for flares related to both Fermi and CTAO, or flares exclusively affecting either Fermi or CTAO. The error bars illustrate the variance of the significance for the 10 randomized realizations of each flare template.

We use scalings of the strength of flares, ranging from 1% to 150% of F_{ref} , given a fixed flare duration of 3 days. The dashed vertical line marks a scaling of 100%. We expect very faint

⁴ <https://github.com/fermi-lat/Fermitools-conda/>

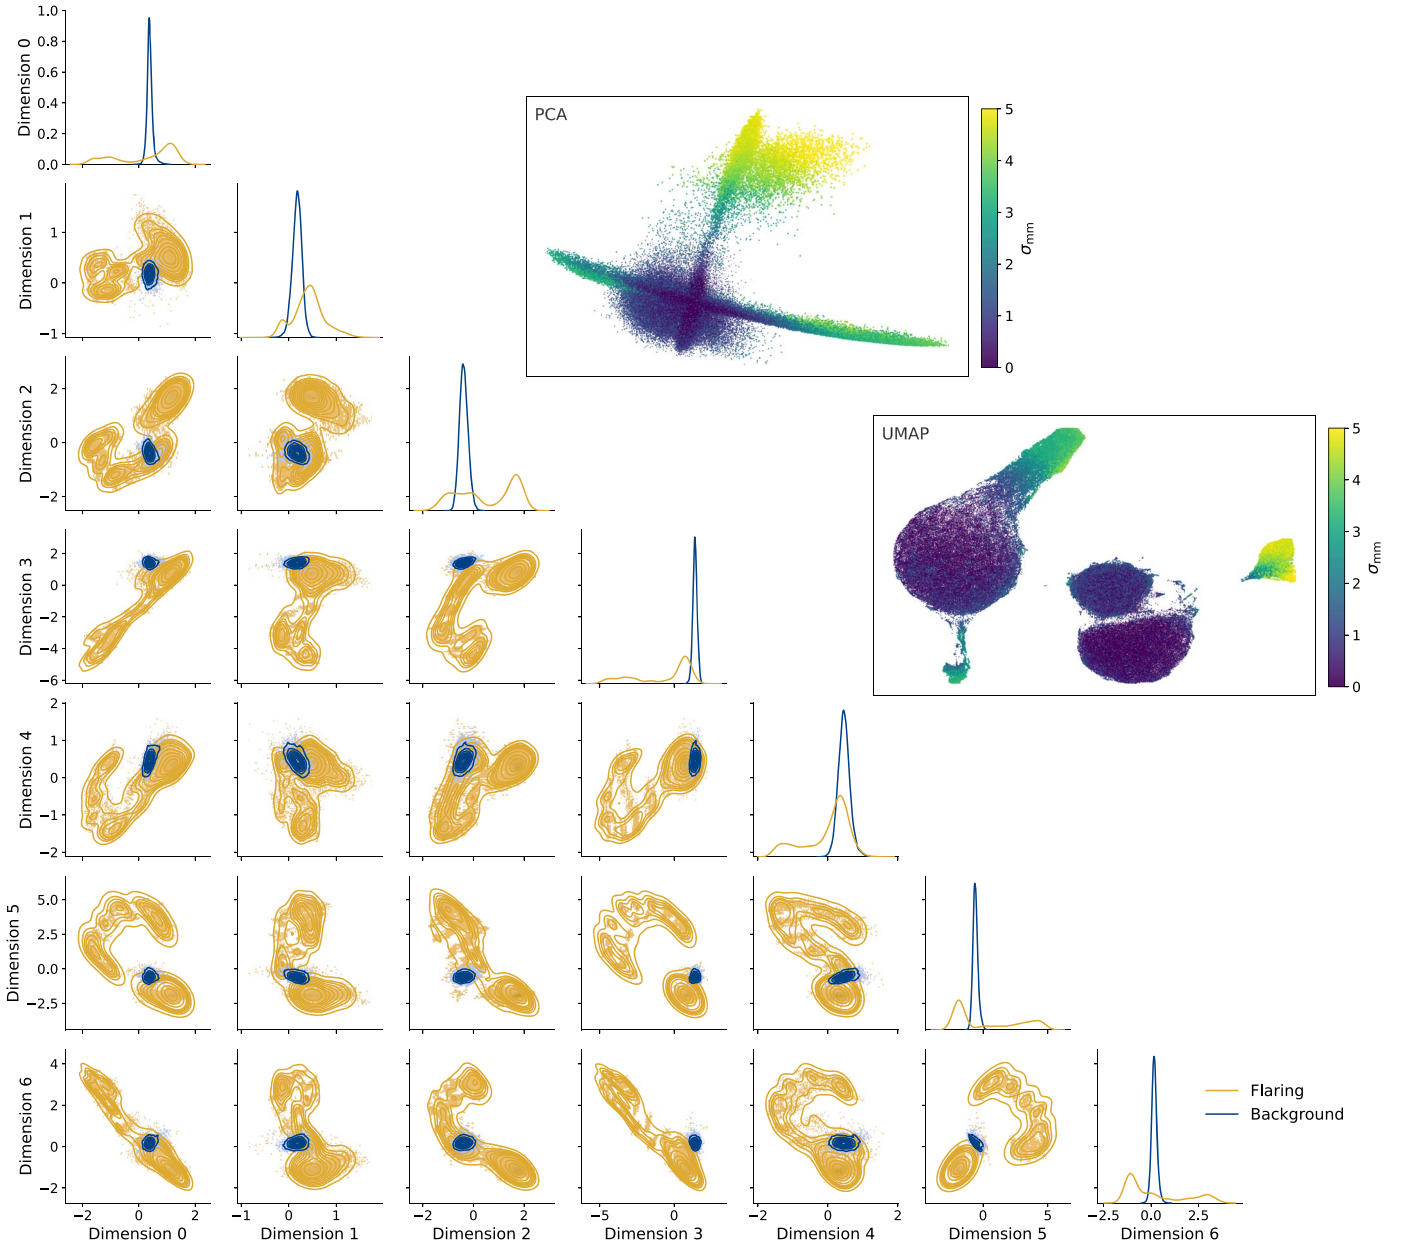


Figure 4. Left: bidimensional projections of the embedding space, showing the distributions of background (blue) and flaring states (yellow). 2D projections are given with PCA on the top left and UMAP on the top right. The color in the PCA and UMAP panels indicates significance (based on the nominal embedding space). For visual clarity, the marker size slightly increases with significance.

flares to be compatible with background states. This is the case, as can be seen by the low significance from flares with a flux of $1\% F_{\text{ref}}$. With increasing strength, the significance rises.

The individual significance values from the mixture model and the reconstruction error are similar; both contribute to the combined significance. For strong flares, i.e., scalings above 100%, the significance from the mixture model appears to saturate. However, the significance of the reconstruction error continues to increase with the strength of flares. Consequently, the increase is also continuous in the combined significance.

The highest significance is reached when a flare is present in both channels. When a flare is present only in CTAO, the corresponding significance is slightly lower. Flares that only occur in Fermi result in significance values lower by a factor of a few. The only exceptions are faint flares ($1\% F_{\text{ref}}$), which always result in significance values close to zero.

It is expected that the significance of Fermi-only flares is lower than CTAO-only flares. This is because Fermi has larger flux uncertainties than CTAO, and therefore the same flare appears less distinct from background variability in Fermi than in CTAO. This is apparent by the fact that we obtain a significance $>5\sigma$ from a CTAO-only flare with 10% strength, whereas a Fermi-only flare with 50% only shows a significance of around 2.5σ . The significance of flares occurring in both channels is generally higher than the significance of a flare in only a single channel.

Besides flaring activity corresponding to elevated short-term states in observables, the framework is also sensitive to short-term decreases in input values. We show this in the right panel of Figure 6. Here, a flare is modeled by a $15\% F_{\text{ref}}$ decrease from background levels, which is clearly detected with significance $>12\sigma$ for CTAO. However, this deviation is




Figure 5. Left: significance of flares, captured in three subsequent days, as a function of the simulated flux of the reference source, 1ES 1215+303. The dashed line corresponds to the originally reported flux, F_{ref} , from the 2017 flare of the source (J. Valverde et al. 2020). Right: significance assuming the baseline flux of the source, F_{ref} , as a function of the number of consecutive observation days. Different combinations of channels, Fermi and CTAO, are considered, as indicated.

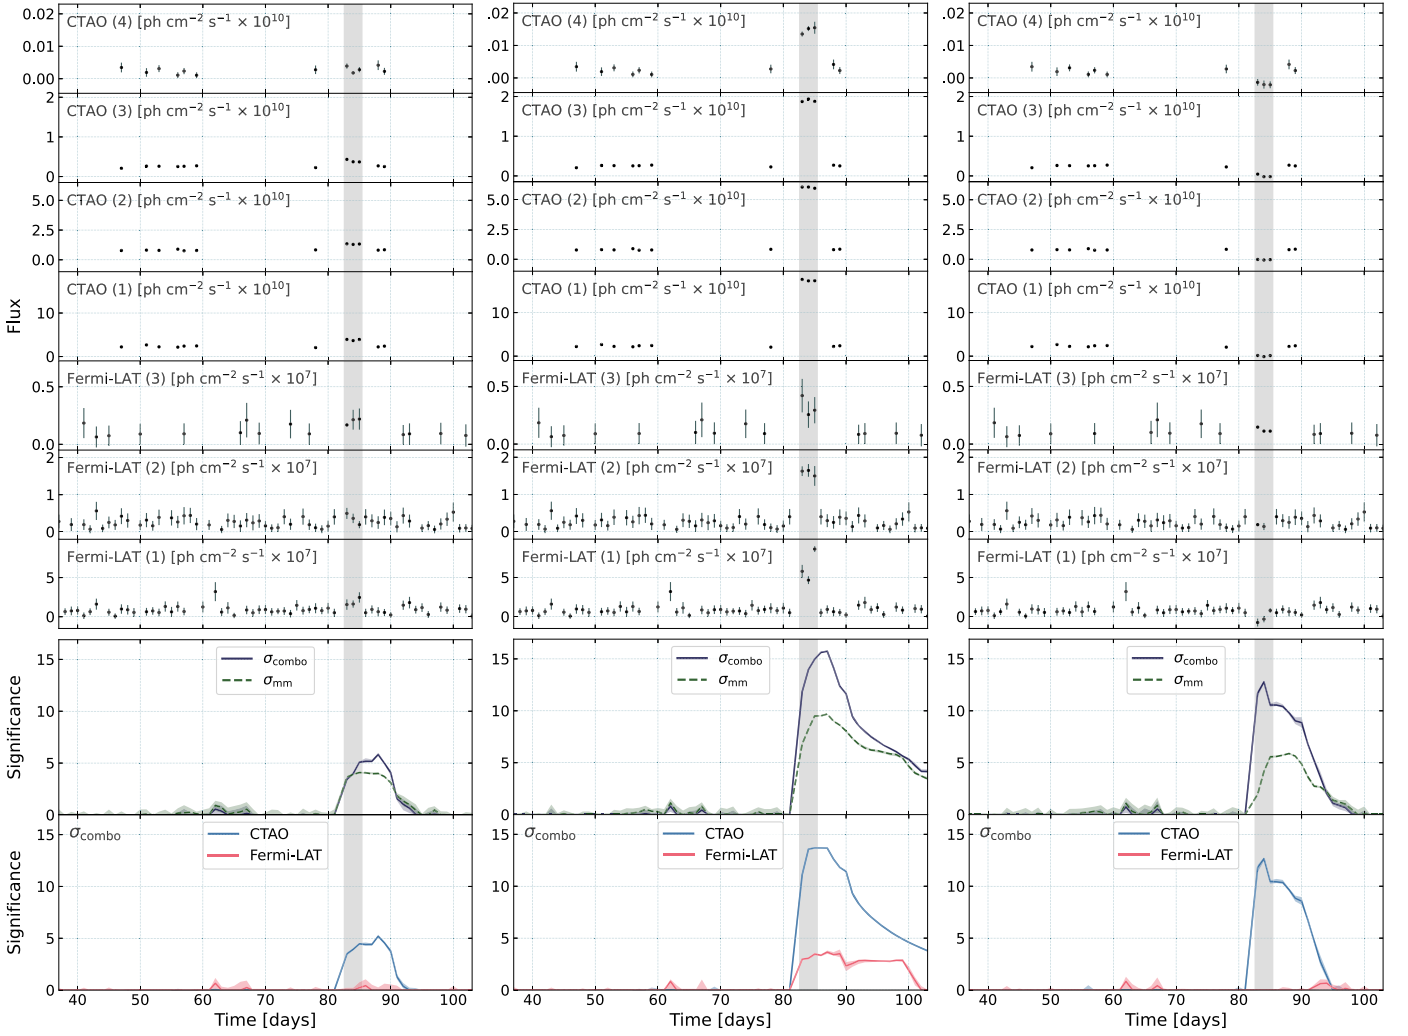


Figure 6. Illustration of injected flare templates, given scaling factors, 10% (left), 100% (middle) and -15% (right), applied to the baseline flux, F_{ref} . The top panels show light curves in the various channels, as indicated. Energy bins for CTAO (panel rows 1–4) and for Fermi (panel rows 5–7) are specified in Section 4.1. The second to last panel shows the clustering significance, σ_{mm} , and the combined significance, σ_{combo} . The lowest panel illustrates the contributions of individual channels to the combined significance. Flaring periods are highlighted as gray shaded regions.

comparable with background fluctuations for Fermi, given the large intrinsic uncertainties in this channel. As expected, the significance increases when considering additional channels that exhibit flaring activity.

4.3.3. Time Coverage

Time coverage describes the number of time steps per channel that are affected by a flare. Coverage can greatly

impact the probability for detection. For weak flares, wider time coverage is required in order to identify anomalies. Conversely, even a single time step might suffice to confidently claim the detection of a strong flare.

We show the impact of time coverage on detection significance in the three panels on the right of Figure 5. Here, we simulate observations of a flare with 100% F_{ref} and time coverage between one and five time steps. We show the combined significance (third column) and the contributions of the individual significance components (last column), also differentiating between detections in only some or in all channels.

Most notably, the same trends are observed related to independent detections in CTAO or in Fermi, as described in the previous section. Flares detected only in Fermi generally result in lower significance. This is again due to the Fermi light curve being intrinsically noisier than the CTAO data. Furthermore, we find that a second detection channel, even if much less significant on its own, can increase the overall significance.

Flares detected in CTAO or in CTAO and Fermi reach a high combined significance with a single high-flux observation. With each additional day of coverage, the significance increases more slowly, and begins plateauing beyond 3 days. Flares only present in Fermi channels, on the other hand, yield a low combined significance at a coverage of 1 day; the significance then increases continuously, reaching a plateau beyond 4 days. This is consistent with the expectation that a weak flare would require longer coverage in order to become significant.

One may observe that the reconstruction significance (bottom-right panel) dominates the overall significance for all cases where a flare is detected in CTAO. This can be explained by the makeup of the training data set: it consists of random, short-term fluctuations with different amplitudes and durations. Due to the intrinsic sparsity in the CTAO channels, only a few cases of fluctuations having coverage of more than one time step appear in the training data. Limited training data for this particular pattern of multiday flares results in potentially greater reconstruction errors. Conversely, the Fermi training data, due to its much denser sampling, contains relatively more multiday fluctuations. The corresponding reconstruction significance is therefore much less dominant. The impact of incomplete training data is mitigated by the complementary use of both σ_{mm} and σ_{rec} .

4.3.4. Time Profile of a Detection

We expect the significance to rise in the presence of new flaring data, and then gradually to fall as physical time passes. We study this behavior in Figure 6. We show three zoomed-in time lines with light curves and their corresponding significance values for a weak upward flare (+10% F_{ref}) on the left, and a strong upward flare (+100% F_{ref}) in the center. On the right, we show a moderate downward flare (−15% F_{ref}). The lowest panel illustrates the contribution to the detection significance when considering individual channels. The flare is highlighted as the gray shaded region.

As expected, the significance starts rising as soon as the first flaring time step comes into play. It reaches a plateau just after the end of the flaring period, and then decays with time. In the case of the weaker flare on the left side of Figure 6, there is a small peak after the flaring period. Here, a time step in the

topmost channel exhibits comparable strength to the preceding flare. The continuously elevated state on this timescale is incompatible with background, resulting in high significance. Thereafter the significance falls, as the variance of the following data returns to background levels. The example of injecting a negative fluctuation illustrates that the framework is agnostic to the topology of potential flares.

5. Historical Blazar Study

After evaluating our framework on simulations, we proceed to analyze historical data and flares of the blazar BL Lacertae (BL Lac). We assess the effectiveness of our framework by comparing with the reactions of the real-time multiwavelength community. In this context, we make use of alert systems, such as the General Coordinates Network (GCN) and the Astronomer’s Telegram (ATel). ATels can be published with a delay of up to a few days. We therefore compare the date the framework detects an anomaly to the time of observation, rather than to the date of publication.

5.1. Data Set

We obtained multiwavelength data of the blazar BL Lac from a number of instruments operating in different parts of the electromagnetic spectrum. The data span a time range of ~ 4000 days, MJD 55170–59200 (between 2009 December and 2020 December). Observation times are rounded to a timescale, $\gamma_{\text{timeaggr}} = 1$ day. A combined multiwavelength light curve for these data is shown in the top panels of Figure 7. A detailed description of the individual data sets and of the related processing and analysis steps is given in Appendix B.

This light curve includes data from the following observatories. In the VHE band between 100 GeV and 30 TeV, we use data from VERITAS, an array of four IACTs in Arizona (J. Holder et al. 2008). While VERITAS has published several distinct emission episodes of BL Lac (T. Arlen et al. 2013; Q. Feng et al. 2017; A. U. Abeysekara et al. 2018), the complete data set remains private. The collaboration has kindly allowed us to explore all available data for our analysis. Accordingly, we use the flux of the source as input to our pipeline. For visual purposes only (Figures 7 and 8 in the following), we mask the data. Instead of flux, we show the corresponding excess photon count rate, slightly modified by random noise. Though the data presented in the figures are in arbitrary units, they track the relevant changes to the measured flux of the source. They can thus be used to appreciate the performance of our method, as discussed below.

For energies between 100 MeV and 300 GeV, we use data taken by the LAT on board the Fermi spacecraft (W. B. Atwood et al. 2009), divided into the same three energy bins as before (100–669 MeV, 669 MeV–4.48 GeV, and 4.48–300 GeV). In the X-ray band, we include observations of BL Lac taken by the X-Ray Telescope (XRT) on board the Swift satellite (D. N. Burrows et al. 2005), divided into soft (0.3–1.5 keV) and hard (1.5–10 keV) bands.

Optical r -band data include observations taken with the Tuorla optical telescope in Finland (K. Nilsson et al. 2018), observations taken with the Palomar Transient Factory (PTF) monitoring program (N. M. Law et al. 2009) in San Diego County, California, and with the successor instrument, the Zwicky Transient Facility (ZTF; E. C. Bellm et al. 2019a, 2019b). The latter two observatories also provide g -band magnitudes of the source.

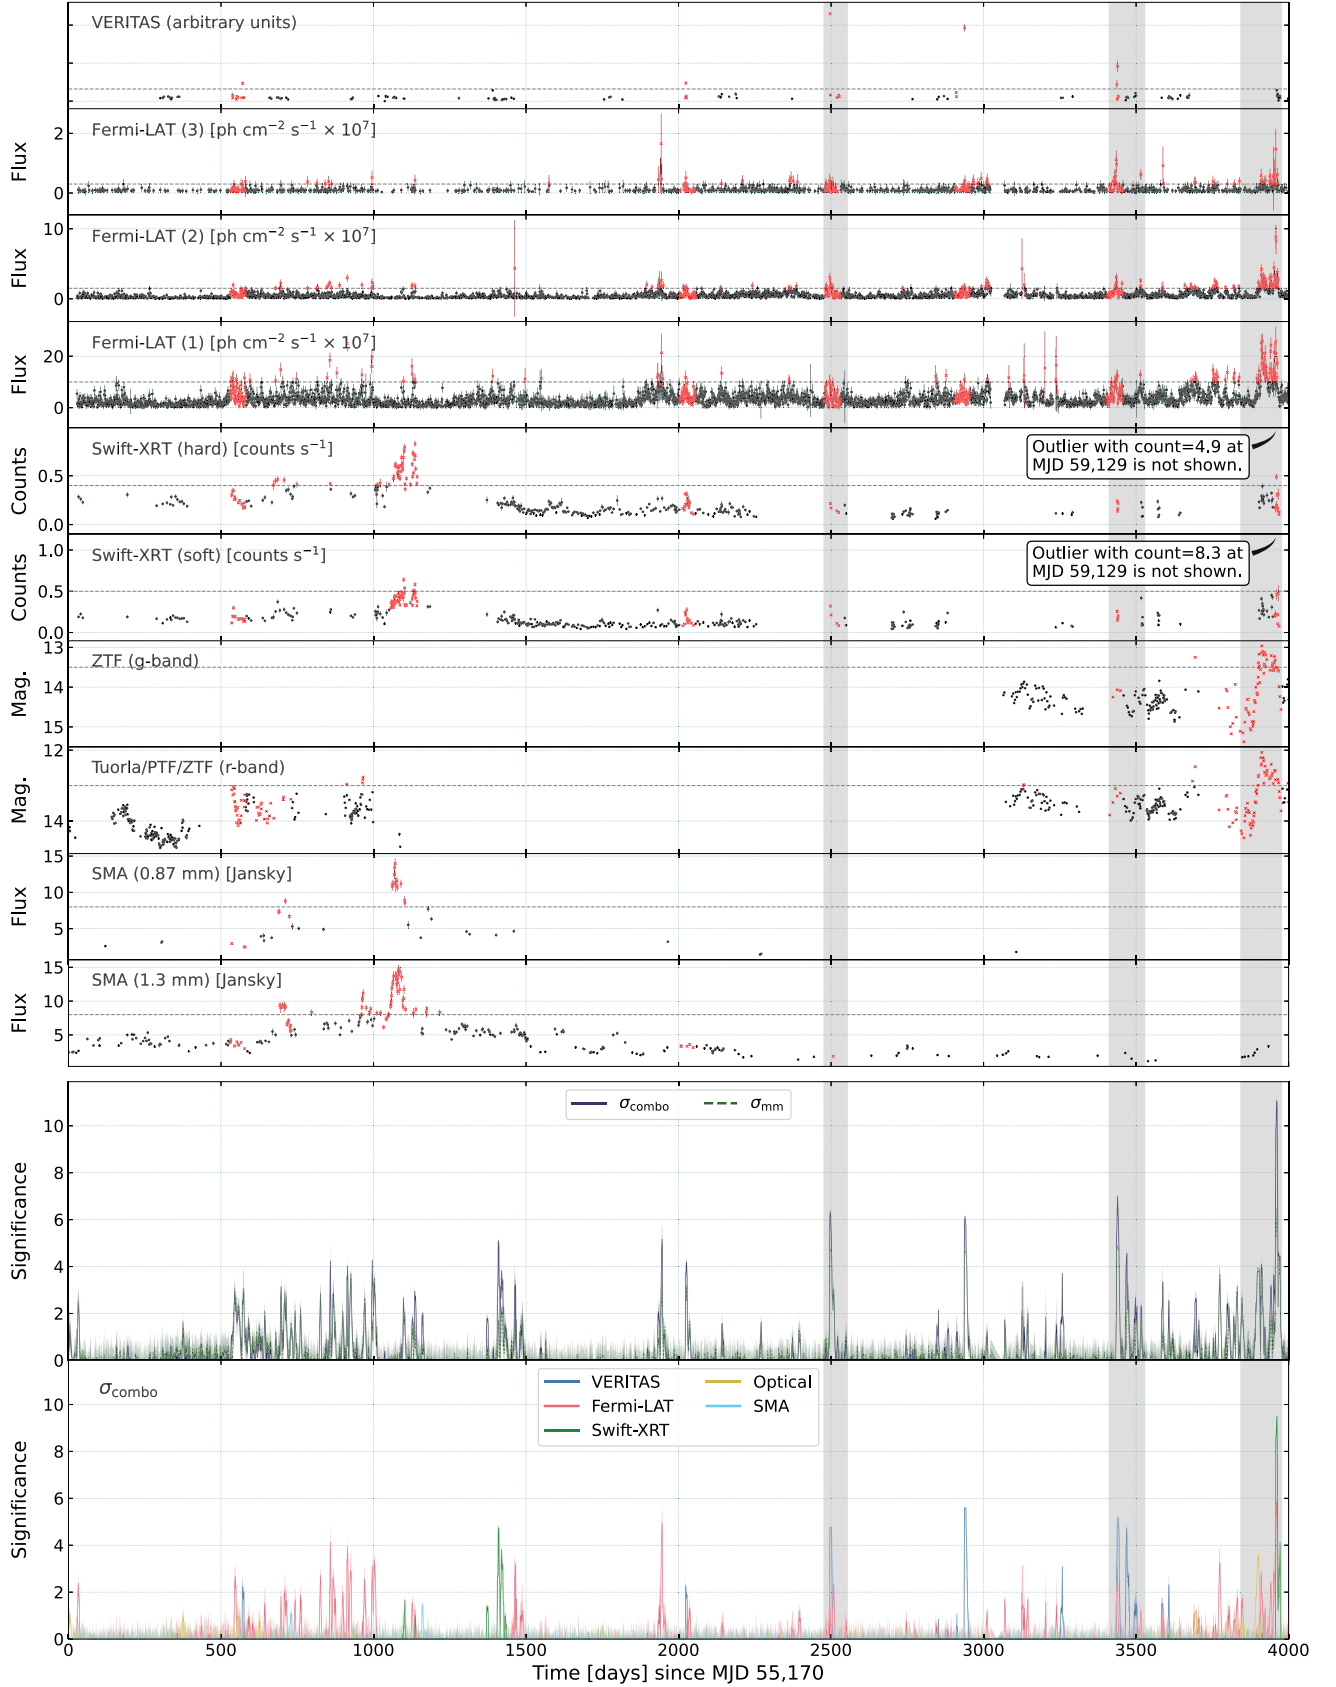


Figure 7. Overview of the multiwavelength BL Lac data set and the corresponding significance of flares, σ_{mm} and σ_{combo} . The lowest panel illustrates the contributions of individual channels to the combined significance. Data marked with red crosses are excluded from the training and calibration samples. Horizontal dashed lines indicate the associated background cuts. Gray shaded regions mark the time ranges highlighted in Figure 8. Both Swift-XRT channels contain an outlier, hidden for readability.

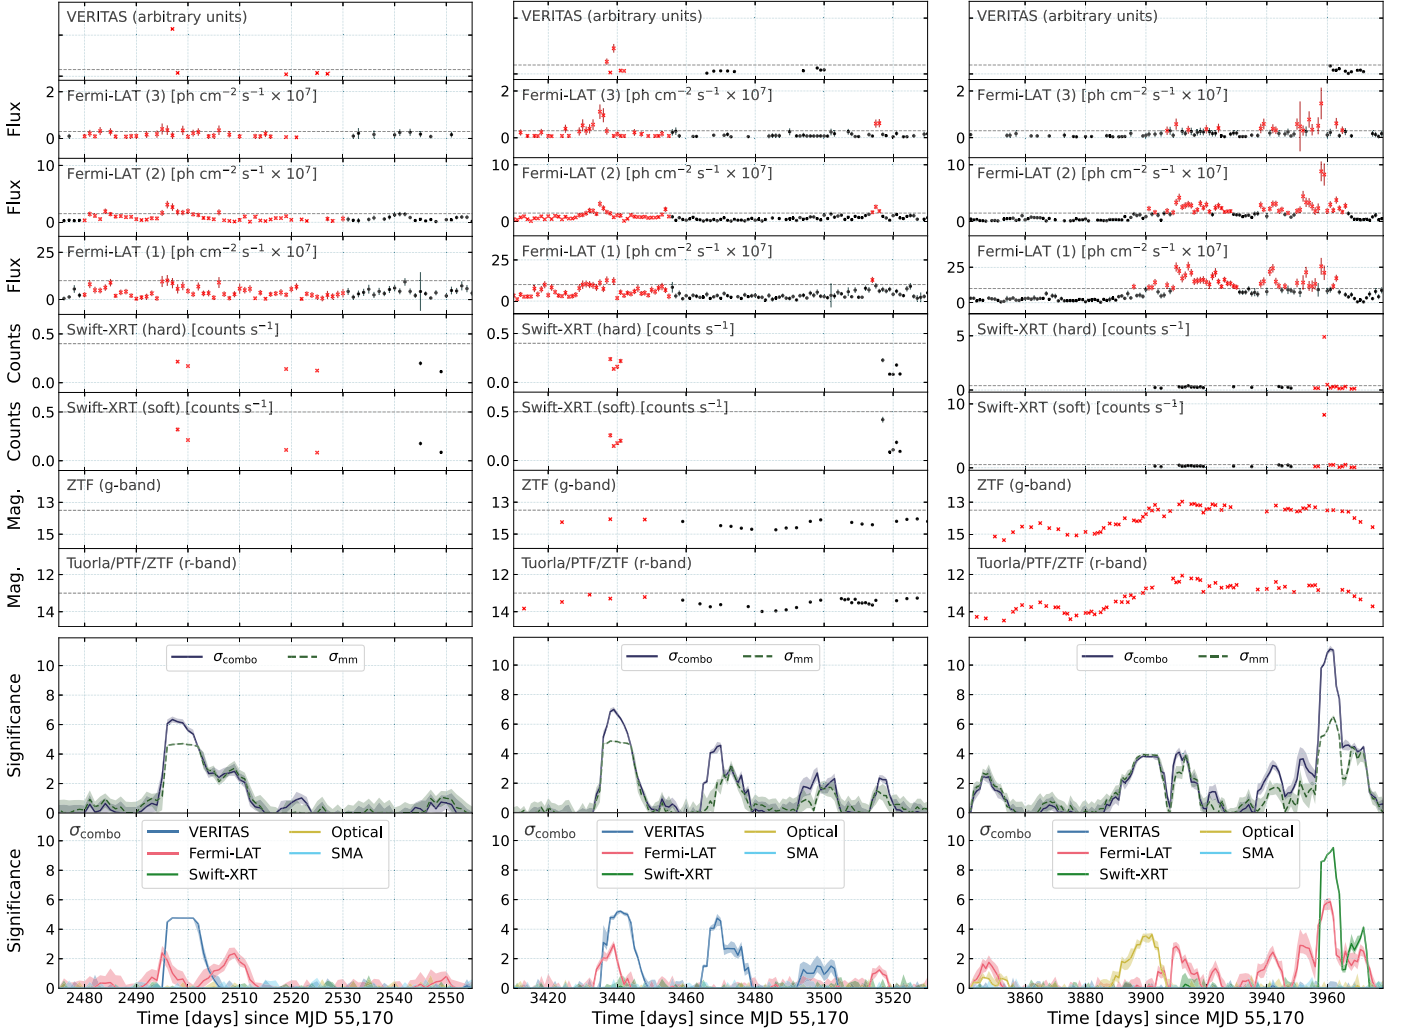


Figure 8. Highlighted intervals related to high-significance detections in the BL Lac data set (see caption of Figure 7).

Finally, we use radio/microwave data at wavelengths of 0.87 mm and 1.3 mm, taken by the Submillimeter Array (SMA) located on Maunakea in Hawaii (M. A. Gurwell et al. 2007).

Data from VERITAS, Fermi-LAT, and SMA are in units of flux, data from Tuorla and ZTF/PTF are in magnitude (Mag.), and data from Swift-XRT consist of count rates. We decided against performing an advanced X-ray analysis, which would include using a spectral model and fitting the neutral hydrogen column density. Instead, we use the lower-level data product, provided through the Swift online analysis tools (see also Appendix B). By design, the framework is agnostic to using different types of input across channels. We approximate the uncertainties on inputs by sampling; we use Gaussian distributions, having standard deviations corresponding to the quoted uncertainties.

In order to create a training sample that exclusively comprises quiescent states, we first exclude data exhibiting exceptional source activity. The definition of exceptional activity as part of data cleaning determines the sensitivity threshold of the model. We decide to first manually exclude clearly identifiable high states from the light curves. In particular, we remove flaring states based on VERITAS data surrounding MJDs 55725, 57200, 57675, 58100, and 58600.

We exclude any data taken within 50 days of the respective VHE flares from all channels.

We proceed to perform algorithmic cleanup. We exclude all data for a particular time period where at least one of the channels exceeds a given cutoff. Cutoff values are defined in terms of the variance of a channel. Specifically, we exclude data above a maximum flux threshold in the VERITAS, Fermi-LAT, and SMA channels, above a maximum count rate in the Swift channels, and below a minimum magnitude for Tuorla and ZTF/PTF.

5.2. Configuration Parameters and Training

The overall model and parameterization of training are analogous to those used for the simulation study. However, we increase the complexity of the model, as well as the regularization parameter. While we do not perform an exhaustive hyperparameter optimization, we manually adjust some key parameters, in order to minimize underfitting and overfitting. Specifically, we double the size of the RNN layers, double the dropout rate to 20%, and increase the L2 norm to 0.001.

5.3. Results

In Figure 7, we show the full multiwavelength light curve of BL Lac used in this work. The outcome of the data cleaning process is indicated, where all excluded time periods are highlighted by red crosses. For context, the automatically derived cleaning cutoffs are shown as horizontal dashed lines. Time intervals highlighted in gray are shown separately in Figure 8. (The SMA channels are hidden for readability, as they do not contribute notably to the significance for detection.) The two panels on the bottom show the combined significance, and the clustering significance for all channels. The individual contributions to the significance from each channel are illustrative, as discussed above.

In the following, we refer to t as the time (number of days) since MJD 55170.

5.3.1. Performance on Realistic Light Curves

The data used here are much more complex than the simulated light curves investigated in Section 4. We utilize a total of 10 channels. These cover a wide range of the electromagnetic spectrum, having differing observing cadences and occasional extended gaps. The 4000 days light curve contains multiple flaring periods across different channels. From the corresponding significance, one can clearly distinguish episodes of high activity from quiescent periods. This shows that the framework is able to distinguish between different blazar states.

The first flare contained in this light curve is a VHE flare, detected by VERITAS at $t = 570$ (R. A. Ong 2011). This flare is fainter than subsequent VHE flares, and is detected by the framework at $\sim 2.5\sigma$. Three of the highest-significance peaks ($t = 2500$, 2900 , and 3400 days) coincide with VERITAS flares, each reaching $\sim 6\sigma$. A close-up view of the first peak is shown in the left panel of Figure 8. The first and third of these flares are accompanied by a flare detected by Fermi-LAT, as indicated by the individual channel contributions to the significance.

In general, we find elevated significance during the parts of the light curve that were excluded for training. The main exception to this is the significance peak at $t \sim 1450$. At first glance, the surrounding light curve does not show increased variability. The signal here is mostly driven by the reconstruction significance. The lower panel reveals that the largest contribution comes from the Swift-XRT channels. This can be explained by the fast transition from a clear high state into a low state. The transition is first observed at $t \sim 1300$, and continues after a gap in Swift coverage.

The SMA data yield low-significance detections with a $\sim 1\sigma$ spike at $t \sim 740$ and a $\lesssim 2\sigma$ spike at $t \sim 1250$, both following fast transitions between activity states in the respective channels. The two peaks preceding the $\sim 2\sigma$ detection at $t = 1250$ do not result in pronounced signals. This is presumably due to the timescale of our analysis being days. The respective evolution of the SMA light curves is slower; it also exhibits relatively high intrinsic variance, which increases the corresponding sensitivity threshold for anomalous behavior.

5.3.2. Uncertainty Propagation

In Figure 9, we illustrate the impact of input uncertainties on the output of our model. We show a zoom in of the Fermi-LAT

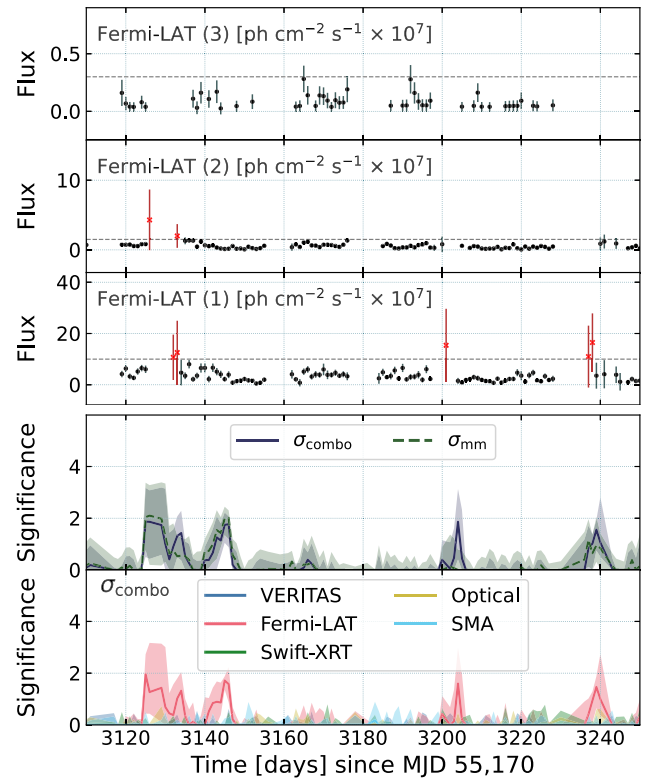


Figure 9. Highlighted interval from the BL Lac data set (see caption of Figure 7). The high significance can be traced back to the Fermi-LAT channel. Large uncertainties on these data yield correspondingly high uncertainties on the detection significance. This can be interpreted as a likely false-positive signal.

channels for $t \sim 3200$. Here, multiple data points have relatively high values, but these are also characterized by large uncertainties. While the corresponding median significance increases to $\sim 2\sigma$, the related uncertainty is quite high. We can thus infer that the elevated state is likely a false-positive signal.

5.3.3. Real-time Detection and Prediction of Increased Multiwavelength Activity

In this section, we compare the performance of the framework with the response of the multiwavelength community to two historical flaring episodes. A summary is provided in Table 2; we show the light curves of these flares in the two panels on the right of Figure 8.

The panel at the center shows a VERITAS flare between $t = 3436$ and 3438 , which is easily identified by the framework at $\sim 5\sigma$. Additionally, the flux in the Fermi channels, most notably in the highest-energy channel, Fermi-LAT 2, starts increasing around $t = 3430$. This precursor is also picked up by the framework, as we can see in the bottom panels. One can observe an increase in significance, driven by Fermi. For the initial flux increase in the Fermi channels, the framework yields a combined significance of $\lesssim 2\sigma$ around $t = 3435$. On $t = 3434$, this trend was noted by the Fermi-LAT collaboration, as reported in an ATel on the following day (S. Garrappa & S. Buson 2019). This ATel was then used to trigger observations with VERITAS and the MAGIC telescope at $t = 3436$; VHE flaring activity was found and subsequently reported by R. Mirzoyan (2019).

Table 2
Summary of the Community and Framework Response to Historical Multiwavelength Flaring Events

t (days)	Event	Framework Response (Significance at Time)	Historical Community Response
570	VHE flare	2.5σ at $t = 570$	Detected at $t = 570$ and reported in ATel at $t = 571$ by VERITAS (R. A. Ong 2011).
2496	VHE flare	6σ at $t = 2496$	Detected at $t = 2496$ and reported in ATel at $t = 2496$ by VERITAS (R. Mukherjee & VERITAS Collaboration 2016).
3434–3435	Flux increase in Fermi channels	1.3σ at $t = 3434$, and 1.7σ at $t = 3435$	Detected at $t = 3434$ and reported in ATel at $t = 3435$ (S. Garrappa & S. Buson 2019).
3436	VHE flaring activity	5σ at $t = 3436$	Detected by VERITAS and MAGIC at $t = 3436$ and reported at $t = 3436$ (R. Mirzoyan 2019).
3470–3500	Return to low state in VHE	4σ at $t = 3467$, and 2.7σ at $t = 3498$	No response.
3880–3900	Brightness increase in optical and gamma-ray flare	3σ at $t = 3896$, and 3.8σ at $t = 3899$	Detected by Crimean Astrophysical Observatory and Fermi-LAT at $t = 3900$ and reported at $t = 3901$ (C. C. Cheung 2020; T. S. Grishina & V. M. Larionov 2020).
3909–3911	Increase in optical and gamma-ray variability	3σ at $t = 3909$	Detected at $t = 3910$ by Fermi-LAT and MAGIC, reported at $t = 3913$ (O. Blanch 2020a; R. Ojha & J. Valverde 2020). Optical brightening detected by the Automatic Telescope for Optical Monitoring between $t = 3909$ and $t = 3911$, reported at $t = 3911$, detected by Hans-Haffner-Sternwarte at $t = 3910$, reported at $t = 3912$ (F. Jankowsky & S. Wagner 2020; R. Steineke et al. 2020).
3940–3941	Gamma-ray flare	2.4σ at $t = 3940$ (driven by Fermi data)	Detected by MAGIC at $t = 3941$ (O. Blanch 2020b)
3952–3957	Increased gamma-ray activity	3.6σ at $t = 3951$, and 5σ at $t = 3957$	Detected by Fermi-LAT collaboration from $t = 3957$ onwards, reported $t = 3959$ (I. Mereu 2020).
3958	Record X-ray flare	10σ at $t = 3958$	X-ray and UV rise detected by Swift at $t = 3957$, reported at $t = 3958$ (F. D’Ammando 2020a), X-ray flare detected by Swift-XRT at $t = 3958$, reported at $t = 3959$ (F. D’Ammando 2020b).

Note. The parameter t denotes the time in days since MJD 55170.

After this flare, subsequent observations with VERITAS around $t = 3470$ and 3500 reveal a return to a low-activity state. For these observations, our framework yields a significance of $\sim 4\sigma$ and $\sim 2.7\sigma$, respectively, driven entirely by the VERITAS channel. This shows that abrupt transitions from high to low states can also produce a significant detection, provided the low state is observed shortly after the high state.

The panel on the right shows a prolonged phase of increasing multiwavelength activity in 2020, culminating in a brief, very bright X-ray flare at $t = 3958$. From $t \sim 3880$ onwards, the optical brightness slowly starts increasing, with the significance beginning to rise roughly 7 days later. The significance reaches 3σ at $t = 3896$, and 3.8σ at $t = 3899$. A day later, the Crimean Astrophysical Observatory observed BL Lac in its historical optical r -band maximum. The Fermi-LAT collaboration observed a gamma-ray flare at the same time, both reported in ATels the following day (C. C. Cheung 2020; T. S. Grishina & V. M. Larionov 2020).

Up to this point, the significance is driven purely by the continuing increase in the optical channels, and the brightening gamma-ray flux is not significantly detected by the framework. However, after a short period of fluctuation in the g -band and Fermi-LAT channels, the framework yields another 3σ detection for $t \sim 3909$. In this case, it is driven by a continued increase in gamma-ray variability. Observations with Fermi-LAT and the VHE gamma-ray telescope, MAGIC, revealed a flare a day later ($t = 3910$). This was subsequently reported in ATels (O. Blanch 2020a; R. Ojha & J. Valverde 2020). No

contemporaneous observations of this MAGIC-detected VHE flare were taken with VERITAS.

Simultaneously, continued optical observations with multiple instruments also found BL Lac in a high state at $t = 3909$ (F. Jankowsky & S. Wagner 2020; R. Steineke et al. 2020). A further increase in Fermi-LAT flux is detected around $t = 3940$ at 2.4σ . At $t = 3941$, the MAGIC collaboration again detected a bright VHE gamma-ray flare (O. Blanch 2020b). Here, the framework yields 2.4σ one day earlier. At $t = 3951$, the continued increase in gamma-ray variability is detected by the framework at 3.6σ .

These increases in multiwavelength activity were used to trigger observations with the Swift satellite for 4 days (3954–3957). The XRT observed the second-highest count rate from BL Lac up to that date, and the Swift UV-Optical Telescope found the source to be in an elevated UV-optical state at the same time (F. D’Ammando 2020a). Further Swift observations at $t = 3958$ subsequently found a “record X-ray flare,” with the highest observed count rate from this source (F. D’Ammando 2020b). Also at $t = 3958$, Fermi-LAT observed a gamma-ray flare from BL Lac (I. Mereu 2020). Meanwhile, the framework reaches a significance around 5σ a day earlier, at $t = 3957$, driven by the high flux in all Fermi channels, particularly the HE ones. This significance then quickly exceeds 10σ with the detection of the exceptionally bright X-ray state at $t = 3958$. Such a high X-ray state is quite divergent from the data used to train the framework. Accordingly, the reconstruction significance contributes strongly to the combined signal.

5.4. Conclusions

We demonstrate the ability of our model to reliably detect flares of the blazar BL Lac, as well as hint at their precursors. Detected events can be further characterized with our model. This is accomplished by retracing the contributions of the different channels (or combinations thereof) to the final significance. We show that the uncertainties on the multi-wavelength observables can be propagated onto the corresponding significance for detection; this improves the reliability predictions.

Our method is suited for real-time analyses. Comparing to the historical response of the community, we show that the framework detects flaring states at the same time or earlier than the state-of-the-art. The significance obtained from our model offers a simple, single-scalar metric to quantify anomalous states (and possible precursors) in multiwavelength light curves. Such a metric can be helpful for communication. Our approach thus has the potential to improve the reaction of the community to interesting events.

6. Summary and Discussion

In this work, we showcase a novel deep learning analysis framework, capable of detecting various types of anomalies in real-world, multiwavelength light curves. While handling varying cadences and gaps in observational coverage, our model differentiates source variability from noisy background activity. As we avoid the need for a labeled training data set of flaring states, we open the door to discovery of the unexpected.

We evaluate our framework on simulations and on historical blazar light curves. We demonstrate that we can effectively detect flaring behavior consisting of clear high or low states, as well as of correlated signals across bands. An advantage of our approach is the standardized quantification of anomalous states in terms of their significance. This can be used to robustly make consistent decisions on resource management for future observations. It also facilitates communication within the multiwavelength community, where different definitions of a flare are commonly used.

Performance on simulations and on real data. Using light-curve simulations, we show that upward and downward fluctuations from the baseline state of a source are detectable. The associated significance scales with the intensity of deviations. Strong fluctuations from the baseline can already be identified from a single data point. Fainter deviations or those in noisy channels require multiple subsequent data points for meaningful detection, as one would expect.

We evaluate real observations of blazar variability spanning a wide range of the electromagnetic spectrum. Our framework can reliably detect known historical flares as they occur in different channels. The respective contributions to the detection significance from different channels (or combinations thereof) can also be traced. Comparing our results to the reaction of the international community, we confirm that our framework can consistently pick up flares on the same timescales. In some cases, our model would have provided hints of upcoming interesting activity earlier than was achieved in practice.

Timescales of source variability. A crucial characteristic of blazars is their pronounced variability on various timescales, ranging from minutes to years. By design, an instance of the framework is sensitive to variability on a specific timescale. (This is adjustable per channel, given the minimum time bin

width parameter, γ_{timeaggr} .) In the current study, we illustrate the sensitivity of our framework for variability timescales of days. Other timescales could be considered in the future, e.g., using an ensemble of trained pipelines, covering timescales from hours to weeks.

The choice of a timescale for the search window is informed by the science questions being investigated. It is also constrained by the availability of relevant data. To resolve this, one could for example augment existing data with simulations. These may span different quiescent configurations, usable as new training samples. Such an approach also opens the door to tuning the sensitivity of the model to particular topologies of flares.

Role of data cleaning. As the framework is trained on nonanomalous data, which rarely exists in practice, prior data cleaning is needed. This may introduce biases, as the definition of anomalous source behavior is typically influenced by the expectations and goals of the analysis. If the cleaning procedure is too lenient, the framework will not be sensitive to subtle anomalies. Conversely, cleaning that is too strict leads to an increase in the number of false-positive signals. For some highly variable sources, identifying and excluding anomalous states can be challenging, especially when there are large gaps in coverage. To address this challenge, our approach incorporates a statistical data cleaning technique based on the γ_{signoise} parameter.

Impact of model dimensionality. In general, simple models may not capture the richness of a data set, and could result in underfitting. Increasing the number of dimensions enables the framework to model more complex source variability. However, this comes at the cost of additional computational complexity, and may result in overfitting.

As part of the forecasting stage, we address possible underfitting by progressively increasing the size of the model. This is moderated by verifying that the training score no longer improves past a certain level of complexity. To avoid overfitting, we use a wide range of augmented and randomized realizations of the training data set, as described in Section 3. We also incorporate standard regularization techniques.

Another crucial aspect of the framework is the complexity of the autoencoder, and in particular the size of the embedding vector. The number of components of the BGMM is also of great importance. As described above, we gradually increase the complexity until the mapping to clusters in the embedding space stops improving. We use a diverse collection of augmented and randomized inputs, as for the forecasting stage. Augmentation in this case also includes addition of upward and downward fluctuations to the data. Such examples reduce the reconstruction error of the autoencoder for potential flares, which are not present in the background training sample. The performance of the autoencoder is also verified as part of inference, by inspection of the reconstruction error via σ_{rec} . Systematically high values of the reconstruction error would indicate underfitting.

Resampling uncertainties on observables. We have illustrated how our model lends itself to propagating uncertainties on input data. In particular, outliers with substantial uncertainties yield significance values with respectively large uncertainties. This allows one to effectively avoid false-positive predictions that are related to low-quality data.

In order to propagate uncertainties from observables to the corresponding significance, the input light curves are

resampled. For the current study, the underlying assumption is that the provided uncertainties on inputs are statistically independent; more complicated correlations may be considered as needed. We also note that after resampling individual values may be nonphysical (e.g., correspond to negative fluxes). While a restriction on the resampling process could circumvent this, that would lead to an underestimation of the corresponding uncertainty. In any case, the data-driven neural network approach is agnostic to such transformations on input data.

Impact of varying observation cadence. We designed our framework to be robust against potentially confounding changes in cadence. Our solution, based on temporal weights, enhances the sensitivity for recent signals in sparse channels. The downside of this approach is that it does not differentiate between recent low activity and lack of data. This may limit the capability of the model to accurately capture intricate temporal patterns within and across channels. More sophisticated weighting schemes may be considered to address specific use cases.

Future work and extensions. Our approach can be utilized to detect and analyze transient phenomena in general; the main constraint is that the signal is detectable via deviations from an established baseline of background activity. Our framework therefore sets the groundwork for a broad range of future work. In particular, the pipeline is suitable for deployment as part of a real-time automated broker system, such as AMPEL (J. Nordin et al. 2019). It can therefore be used to quickly detect interesting activity and alert the community.

For example, future work may explore more sophisticated forecasting models. In particular, probabilistic modeling of the differences between predictions and real data could increase the robustness of the model.

Furthermore, the framework could be extended to differentiate between known categories of activity, while remaining sensitive to unexpected anomalies. For instance, separate significance metrics could be derived for specific types of variability, e.g., based on clustering in the embedding space. This also has potential as a tool for comparing different theoretical emission models, or for making predictions and detecting flares in the context of specific models.

Acknowledgments

We would like to thank the following people for numerous useful discussions: Q. Feng, M. Gurwell, C. McGrath, M. Negro, D. Parsons, and B. Rani. We would also like to thank the CTAO consortium and the Fermi-LAT and VERITAS collaborations for conducting internal courtesy reviews, and for providing useful feedback on this work.

This work is supported by the Helmholtz Einstein International Berlin Research School in Data Science (HEIBRiDS). This research was supported by the Helmholtz Weizmann Research School on Multimessenger Astronomy, funded through the Initiative and Networking Fund of the Helmholtz Association, DESY, the Weizmann Institute, the Humboldt University of Berlin, and the University of Potsdam.

This work made use of data supplied by the UK Swift Science Data Centre at the University of Leicester.

This research made use of the NASA/IPAC Infrared Science Archive, which is funded by the National Aeronautics and Space Administration and operated by the California Institute of Technology.

This research made use of observations from the Submillimeter Array, a joint project between the Smithsonian Astrophysical Observatory and the Academia Sinica Institute of Astronomy and Astrophysics, funded by the Smithsonian Institution and the Academia Sinica. We recognize that Maunakea is a culturally important site for the indigenous Hawaiian people; we are privileged to study the cosmos from its summit.

This research has made use of the VizieR catalog access tool, CDS, Strasbourg, France (F. Ochsenbein 1996). The original description of the VizieR service was published in F. Ochsenbein et al. (2000).

This research made use of observations obtained with the 48-inch Samuel Oschin Telescope and the 60 inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project. ZTF is supported by the National Science Foundation under grant Nos. AST-1440341 and AST-2034437, and a collaboration including current partners Caltech, IPAC, the Oskar Klein Center at Stockholm University, the University of Maryland, University of California, Berkeley, the University of Wisconsin at Milwaukee, University of Warwick, Ruhr University, Cornell University, Northwestern University, and Drexel University. Operations are conducted by COO, IPAC, and UW.

The Fermi-LAT Collaboration acknowledges generous ongoing support from a number of agencies and institutes that have supported both the development and the operation of the LAT as well as scientific data analysis. These include the National Aeronautics and Space Administration and the Department of Energy in the United States, the Commissariat à l’Energie Atomique and the Centre National de la Recherche Scientifique/Institut National de Physique Nucléaire et de Physique des Particules in France, the Agenzia Spaziale Italiana and the Istituto Nazionale di Fisica Nucleare in Italy, the Ministry of Education, Culture, Sports, Science and Technology (MEXT), High Energy Accelerator Research Organization (KEK) and Japan Aerospace Exploration Agency (JAXA) in Japan, and the K. A. Wallenberg Foundation, the Swedish Research Council and the Swedish National Space Board in Sweden. Additional support for science analysis during the operations phase is gratefully acknowledged from the Istituto Nazionale di Astrofisica in Italy and the Centre National d’Études Spatiales in France. This work was performed in part under DOE Contract DE-AC02-76SF00515.

The authors thank the VERITAS collaboration for making available the data used in this publication.

This research made use of the CTAO instrument response functions; see <https://www.cta-observatory.org/science/cta-performance/> for more details.

Data Availability

The code used in this work is available on GitLab under a MIT license: https://gitlab.desy.de/trans_finder/blazar_flares. The simulation (HDF5) data for this study have been deposited on Zenodo under a Creative Commons Attribution (v4) license: doi:10.5281/zenodo.14698916.

Facilities: Fermi, SMA, Swift, PO:1.2m, PO:1.5m, IRSA, VERITAS.

Appendix A Light-curve Simulations

This section describes in more detail the specifics of the light-curve simulations of 1ES 1215+303. Three representative examples of simulated multiwavelength light curves are shown in Figure 6.

A.1. Fermi-LAT

The Fermitools package (version 2.2.0), a software package provided by the Fermi-LAT collaboration to facilitate analysis of telescope data, includes the **gtobssim** tool for simulating LAT observations.⁵ The inputs to these simulations are the instrument response functions (IRFs) to be used, the time range of the observations, a list of sources to be simulated and parameters corresponding to each source, such as their positions in the sky, their average fluxes, and their spectral models. As the goal is to train the model on pure background data, we simulate the source at a constant flux level, only perturbed by Poissonian- and instrument-induced noise. Because we are simulating the source in two adjacent energy bands, Fermi-LAT and CTAO, respectively, we make sure that the simulated spectra in both bands are connected. In order to achieve this, we use the “SpectralTransient” class from **gtobssim**, allowing us to simulate a log-parabolic spectrum. As input to the simulation, we use the fluxes obtained for the 2017 nonflaring period defined by J. Valverde et al. (2020, Table 9). We simulate a total of 1000 consecutive days of Fermi-LAT observations between MJD 55999 and 56999. In our case, we use the latest LAT IRF (P8R3_SOURCE_V3), and choose to limit the simulation to the source in question. We use the latest templates for diffuse and Galactic background components (gll_iem_v07.fits, iso_P8R3_SOURCE_V3_v1.txt). Additionally, we use the relevant spacecraft file for the time range of the simulations, obtained online.⁶

After generating the simulated event files, we use **fermipy** (v. 1.2, M. Wood et al. 2017) to perform a Fermi-LAT maximum-likelihood analysis, fitting the target source and the background components in three adjoining energy bins (100–669 MeV, 669 MeV–4.48 GeV, and 4.48–300 GeV) over the full 1000 days. We use a 15° circular ROI and apply standard event-quality cuts, (*DATA_QUAL*>0)&& (*LAT_CONFIG* == 1), *evclass* = 128, *evtype* = 3, *zmax* = 90.

Finally, a daily-binned light curve is produced by fitting the source parameters together with the normalization of the isotropic diffuse background for each 1 day interval within each energy bin. For all input light curves, low-confidence data are discarded beforehand. This corresponds to all data with a negative detection TS, and all data where the relative flux uncertainty, $\Delta\text{Flux}/\text{Flux}$, is larger than 5.

A.2. CTAO

Simultaneously, we simulate light curves of 1ES 1215+303 as it would be observed by the upcoming CTAO, using the **gammapy** package (v. 0.20.1; C. Deil et al. 2017) together with the public preliminary CTAO IRF.⁷ Specifically, the IRF we use is Prod5-North-20deg-AverageAz-4LSTs09MSTs.18000s-v0.1.fits.gz. In order to simulate a CTAO light curve, one has to

supply the following inputs: the source coordinates, an assumed spectral model, including average flux and a spectral index, a time profile, and times and exposure durations for the individual observation runs. We use the “ConstantTemporalModel” class to simulate a realistic IACT observing cadence, given actual observation times from the VERITAS data set of the source in question. This allows us to keep a realistic mix of small and large gaps between individual observations.

As a proxy for the source state, we use the flux and spectral parameters of the 2017 quiescent state, as given in J. Valverde et al. (2020). The spectral model the simulations are based upon is the same log-parabola model as used for the Fermi-LAT simulations. We add to this a term accounting for EBL extinction (J. D. Finke et al. 2010), which results in a sharper drop-off of the flux at VHE energies. The subsequent analysis of the simulated data is also performed in **gammapy**. We divide the simulated data in multiple adjacent energy bins (31.6–79.4 GeV, 79.4–199.5 GeV, 199.5 GeV–1.2598 TeV, and 1.2598–12.589 TeV). We discard data with a relative flux uncertainty higher than 3.

Appendix B Multiwavelength Data

B.1. VERITAS

VERITAS is an array of four IACTs, located at the Fred Lawrence Whipple Observatory outside Tucson, Arizona (31.7° N, 110.0° W). The VERITAS instrument and performance are described in detail in J. Holder et al. (2008) and N. Park (2015). Of particular relevance for this study is the VERITAS sensitive energy range, which extends from ~100 GeV to 30 TeV. VERITAS has observed BL Lac since 2008. We obtain a long-term VERITAS light curve of the source, spanning the years 2008–2020 and including 80 hr of quality-selected data (live time). Data were processed via the standard VERITAS calibration and reconstruction pipelines (G. Maier & J. Holder 2017). Data with relative uncertainty higher than 3 were discarded.

As mentioned above, we use the reconstructed flux of the source as input for our pipeline. The full data set for BL Lac has not yet been published by VERITAS. We therefore show the corresponding excess photon count rate, slightly modified by random noise, in Figures 7 and 8. This count rate traces the relevant changes to the measured flux of the source, and is thus shown here as a proxy for the flux of the source.

B.2. Fermi-LAT

The Fermi-LAT data are retrieved from the public LAT database. In order to create a long-term light curve, we analyze the LAT data in the same three energy bins as for the simulation study, using identical event cuts. For each energy bin, we perform an ROI fit of all photons within a 15° circle centered on BL Lac. We take into account all sources from 4FGL-DR3 (S. Abdollahi et al. 2022) within a 20° circle. This ensures that photons leaking from these sources into the ROI are adequately modeled. We produce a daily-binned light curve for each energy bin over the full duration, including all data with nonnegative TS values. Data with relative uncertainty higher than 3 are discarded.

⁵ <https://github.com/fermi-lat/Fermitools-conda/>

⁶ <https://fermi.gsfc.nasa.gov/cgi-bin/ssc/LAT/LATDataQuery.cgi>

⁷ <https://www.cta-observatory.org/science/cta-performance/>

B.3. *Swift-XRT*

The XRT on board the Swift satellite (N. Gehrels et al. 2004; D. N. Burrows et al. 2005) observes the X-ray sky at energies between 0.2 and 10 keV. We select all observations of BL Lac in the relevant time range and use the Swift online analysis tool to automatically create a count-rate light curve (P. A. Evans et al. 2007).⁸ In order to have some information about the spectral evolution of the source, we make use of the fact that the provided light curve can be separated into two energy bins: the soft band, from 0.3 to 1.5 keV, and the hard band, from 1.5 to 10 keV. In each bin, we filter data having a relative uncertainty ≥ 0.25 . We additionally filter data that could be affected by pileup due to high count rates, particularly in the photon-counting mode (>0.5 counts s^{-1} ; see, e.g., J. Ballet 1999).

B.4. *Optical Data Sets*

B.4.1. *Tuorla*

The Tuorla observatory, with its 1 m optical telescope located in Finland, provides *r*-band magnitudes for a collection of TeV blazars (K. Nilsson et al. 2018) between 2002 and 2012. From this database, we obtain the *r*-band light curve of the blazar BL Lacertae for the relevant time range.

B.4.2. *PTF and ZTF*

The PTF was an optical monitoring program utilizing a collection of telescopes located in San Diego County, California between 2009 and 2012 (N. M. Law et al. 2009). In 2017, this was replaced by the ZTF (E. C. Bellm et al. 2019a, 2019b). Both programs are focused on observing transients, but due to their monitoring method also provide long-term light curves of persistent sources such as blazars. Data sets for both programs are publicly available. We make use of their BL Lac *r*- and *g*-band light curves, provided through the IRSA time-series tool (IRSA 2022). PTF and ZTF data sets have to be retrieved individually, and individual cleaning steps have to be taken. For PTF, we include data with *goodflag* = 1, while for ZTF we filter data with *catflag* ≥ 32768 .

B.4.3. *Combining the Data Sets*

We join the three optical data sets into two light curves, one for the *r* band and one for the *g* band. ZTF supersedes PTF and only began operations 5 yr after the end of the Tuorla program. All optical data after MJD 59784, following a long gap, are therefore exclusively ZTF data.

Some temporal overlap exists between the PTF and Tuorla data sets. Even though the photometry between the two telescopes differs, quasi-simultaneous observations from the instruments provide consistent magnitudes. We therefore combine the respective *r*-band data.

ZTF and PTF also use different, though very similar filters and photometric corrections. In the case of PTF, the Sloan Digital Sky Survey is used as a reference, while the Pan-STARRS1 Survey is used for ZTF. Given the large gap between the latest PTF data and the first ZTF data, even a slight offset between calibrations would not impact our results.

As some PTF data are associated with very large uncertainties, we exclude all PTF data that have an absolute

uncertainty larger than the largest uncertainty in the Tuorla and ZTF data sets. This leaves no *g*-band data before the start of operations of ZTF.

B.5. *SMA*

The 0.87 and 1.3 mm flux density data were obtained at the SMA near the summit of Maunakea (Hawaii). They were kindly provided to us by M. Gurwell. We filter all data that have a relative measurement uncertainty ≥ 0.1 . BL Lac is included in an ongoing monitoring program at the SMA. It aims to determine the fluxes of compact extragalactic radio sources, which can be used as calibrators at millimeter wavelengths (M. A. Gurwell et al. 2007). Observations of available potential calibrators are from time to time observed for 3 to 5 minutes. The measured source signal strength is calibrated against known standard objects, typically solar system objects (Titan, Uranus, Neptune, or Callisto). Data from this program are updated regularly, and are available at the SMA website.⁹

ORCID iDs

Hermann Stolte  <https://orcid.org/0000-0003-0047-5842>
Jonas Sinapius  <https://orcid.org/0009-0004-8608-0853>
Iftach Sadeh  <https://orcid.org/0000-0003-1387-8915>
Elisa Pueschel  <https://orcid.org/0000-0002-0529-1973>
Matthias Weidlich  <https://orcid.org/0000-0003-3325-7227>
David Berge  <https://orcid.org/0000-0002-2918-1824>

References

- Aartsen, M. G., Ackermann, M., Adams, J., et al. 2018, *Sci*, 361, aat1378
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *PhRvL*, 119, 161101
Abdollahi, S., Acero, F., Baldini, L., et al. 2022, *ApJS*, 260, 53
Abeyssekara, A. U., Archambault, S., Archer, A., et al. 2015, *ApJL*, 815, L22
Abeyssekara, A. U., Benbow, W., Bird, R., et al. 2018, *ApJ*, 856, 95
Acciari, V. A., Aliu, E., Arlen, T., et al. 2009, *ApJL*, 693, L104
Acharya, B. S., Actis, M., Aghajani, T., et al. 2013, *Aph*, 43, 3
Aharonian, F., Ait Benkhali, F., Angüner, E., et al. 2022, *Sci*, 376, abn0567
Aharonian, F., Akhperjanian, A. G., Bazer-Bachi, A. R., et al. 2006, *A&A*, 457, 899
Aharonian, F., Akhperjanian, A. G., Bazer-Bachi, A. R., et al. 2007, *ApJL*, 664, L71
Albert, J., Aliu, E., Anderhub, H., et al. 2008, *PhLB*, 668, 253
Aleksić, J., Ansoldi, S., Antonelli, L. A., et al. 2016, *Aph*, 72, 76
Arlen, T., Aune, T., Beilicke, M., et al. 2013, *ApJ*, 762, 92
Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, *ApJ*, 697, 1071
Ballet, J. 1999, *A&AS*, 135, 371
Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019a, *PASP*, 131, 068003
Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019b, *PASP*, 131, 018002
Biteau, J., & Meyer, M. 2022, *Galax*, 10, 39
Blanch, O. 2020a, *ATel*, 13963, 1
Blanch, O. 2020b, *ATel*, 14032, 1
Blei, D. M., & Jordan, M. I. 2006, *BayAn*, 1, 121
Burrows, D. N., Hill, J. E., Nousek, J. A., et al. 2005, *SSRv*, 120, 165
Cerruti, M. 2020, *Galax*, 8, 72
Cherenkov Telescope Array Consortium, Acharya, B. S., Agudo, I., et al. 2019, Science with the Cherenkov Telescope Array (Singapore: World Scientific),
Cheung, C. C. 2020, *ATel*, 13933, 1
D'Ammando, F. 2020a, *ATel*, 14065, 1
D'Ammando, F. 2020b, *ATel*, 14069, 1
Deil, C., Zanin, R., Lefaucheur, J., et al. 2017, *ICRC (Bexco, Busan)*, 301, 766
Evans, P. A., Beardmore, A. P., Page, K. L., et al. 2007, *A&A*, 469, 379
Feng, Q., & Lin, T. T. Y. 2016, in IAU Symp. 325, *Astrophysics*, ed. M. Brescia et al. (Cambridge: Cambridge Univ. Press), 173
Feng, Q., VERITAS Collaboration, Jorstad, S. G., et al. 2017, *ICRC (Bexco, Busan)*, 301, 648

⁸ https://www.swift.ac.uk/user_objects/

⁹ sma1.sma.hawaii.edu/callist/callist.html

- Finke, J. D., Razzaque, S., & Dermer, C. D. 2010, *ApJ*, **712**, 238
- Garrappa, S., & Buson, S. 2019, *ATel*, **12718**, 1
- Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, *ApJ*, **611**, 1005
- Grishina, T. S., & Larionov, V. M. 2020, *ATel*, **13930**, 1
- Gurwell, M. A., Peck, A. B., Hostler, S. R., Darrah, M. R., & Katz, C. A. 2007, in *ASP Conf. Ser. 375, From Z-Machines to ALMA: (Sub)Millimeter Spectroscopy of Galaxies*, ed. A. J. Baker et al. (San Francisco, CA: ASP), 234
- H. E. S. S. Collaboration, Abdalla, H., et al. 2019, *Natur*, **575**, 464
- Holder, J., Acciari, V. A., Aliu, E., et al. 2008, in *AIP Conf. Proc. 1085, High Energy Gamma-Ray Astronomy*, ed. F. A. Aharonian, W. Hofmann, & F. M. Rieger (Melville, NY: AIP), 657
- IRSA 2022, Time Series Tool, IPAC, doi:10.26131/IRSA538
- Jankowsky, F., & Wagner, S. 2020, *ATel*, **13956**, 1
- Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
- Kirk, J. G., Guthmann, A. W., Gallant, Y. A., & Achterberg, A. 2000, *ApJ*, **542**, 235
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *PASP*, **121**, 1395
- MAGIC Collaboration, Acciari, V. A., & Ansoldi, S. 2019, *Natur*, **575**, 455
- Maier, G., & Holder, J. 2017, *ICRC (Bexco, Busan)*, **35**, 747
- Malanchev, K. L., Pruzhinskaya, M. V., Korolev, V. S., et al. 2021, *MNRAS*, **502**, 5147
- McInnes, L., Healy, J., Melville, J., et al. 2018, arXiv:1802.03426
- McLachlan, G. J., & Peel, D. 2000, *Finite Mixture Models*, Vol. 299 (New York: Wiley)
- Mereu, I. 2020, *ATel*, **14072**, 1
- Min, E., Guo, X., Liu, Q., et al. 2018, *IEEEA*, **6**, 39501
- Mirzoyan, R. 2019, *ATel*, **12724**, 1
- Mukherjee, R. & VERITAS Collaboration 2016, *ATel*, **9599**, 1
- Murase, K., Dermer, C. D., Takami, H., & Migliori, G. 2012, *ApJ*, **749**, 63
- Nieto Castaño, D., Brill, A., Kim, B., & Humensky, T. B. 2017, *ICRC (Bexco, Busan)*, **301**, 809
- Nilsson, K., Lindfors, E., Takalo, L. O., et al. 2018, *A&A*, **620**, A185
- Nordin, J., Brinnet, V., van Santen, J., et al. 2019, *A&A*, **631**, A147
- Ochsenbein, F. 1996, *The VizieR Database of Astronomical Catalogues*, CDS, doi:10.26093/CDS/VIZIER
- Ochsenbein, F., Bauer, P., & Marcout, J. 2000, *A&AS*, **143**, 23
- Ojha, R., & Valverde, J. 2020, *ATel*, **13964**, 1
- Ong, R. A. 2011, *ATel*, **3459**, 1
- Park, N. 2015, *ICRC (The Hague)*, **34**, 771
- Pearson, K. 1901, *PMag*, **2**, 559
- Reis, I., Poznanski, D., Baron, D., Zasowski, G., & Shahaf, S. 2018, *MNRAS*, **476**, 2117
- Rodrigues, X., Fedynitch, A., Gao, S., Boncioli, D., & Winter, W. 2018, *ApJ*, **854**, 54
- Sadeh, I. 2020, *ApJL*, **894**, L25
- Sadeh, I., Abdalla, F. B., & Lahav, O. 2016, *PASP*, **128**, 104502
- Sironi, L., Petropoulou, M., & Giannios, D. 2015, *MNRAS*, **450**, 183
- Steineke, R., Waller, L., Reinhart, D., et al. 2020, *ATel*, **13958**, 1
- Valverde, J., Horan, D., Bernard, D., et al. 2020, *ApJ*, **891**, 170
- Wood, M., Caputo, R., Charles, E., et al. 2017, *ICRC (Bexco, Busan)*, **301**, 824
- Zimmerman, R., van Dyk, D. A., Kashyap, V. L., & Siemiginowska, A. 2024, *MNRAS*, **534**, 2142